AD-A086 716 APPLIED DECISION ANALYSIS INC MENLO PARK CA F/G 17/2
ANALYTIC PROCEDURES FOR DESIGNING AND EVALUATING DECISION AIDS.(U)
APR 80 A C MILLER, P A MORRIS, R D SMALLWOOD N00014-78-C-0240
UNCLASSIFIED ADA-2017 NL

1 of 2
AD
A086716

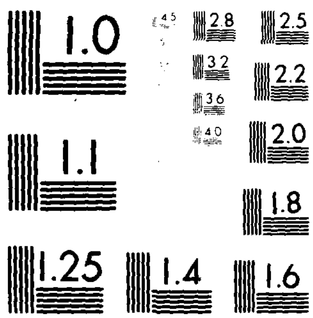

MICROCOPY RESOLUTION TEST CHART

NATIONAL BUREAU OF STANDARDS-1963-A

LEVEL II

12

ADA

APPLIED DECISION ANALYSIS, Inc.
3000 Sand Hill Road, Menlo Park, CA 94025 (415) 854-7101

April 1930

FINAL REPORT

ANALYTIC PROCEDURES FOR DESIGNING AND EVALUATING DECISION AIDS

Prepared for:

Dr. Martin A. Tolcott
Code 455
Director, Engineering Psychology Program
Office of Naval Research
Arlington, Virginia

Prepared by:

Allen C. Miller
Peter A. Morris
Richard D. Smallwood
Robert S. Gibbons

Unclassified

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE

READ INSTRUCTIONS BEFORE COMPLETING FORM

1. REPORT NUMBER / 2. GOVT ACCESSION NO. AD-A086 716	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) (6) Analytic Procedures For Designing and Evaluating Decision Aids.	5. TYPE OF REPORT & PERIOD COVERED (9) Final Report, Apr 79—Mar 80
	6. PERFORMING ORG. REPORT NUMBER ADA Project #2017
7. AUTHOR(s) (10) Allen C. Miller, III; Ronald K. Langford; Peter A. Morris; Robert S. Gibbons; Richard D. Smallwood	8. CONTRACT OR GRANT NUMBER(s) (15) N00014-78-C-0240
9. PERFORMING ORGANIZATION NAME AND ADDRESS Applied Decision Analysis, Inc. 3000 Sand Hill Road Menlo Park, CA 94025	10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS (12) 116
11. CONTROLLING OFFICE NAME AND ADDRESS Office of Naval Research Department of the Navy Arlington, VA 22217	12. REPORT DATE (11) April 1980
	13. NUMBER OF PAGES 105
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office) Defense Contract Administration Services Management Area-San Francisco San Bruno, California 94066	15. SECURITY CLASS. (of this report) Unclassified
	15a. DECLASSIFICATION/DOWNGRADING SCHEDULE

16. DISTRIBUTION STATEMENT (of this Report) (14) ADA-2017

Unlimited

17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)

Unlimited

18. SUPPLEMENTARY NOTES

Reproduction in whole or in part is permitted for any purpose of the United States Government.

19. KEY WORDS (Continue on reverse side if necessary and identify by block number)

Automated Decision Aids
Decision Characteristics
Tactical Command and Control Decisions
Decision Taxonomy

20. ABSTRACT (Continue on reverse side if necessary and identify by block number)

This report describes research conducted by Applied Decision Analysis (ADA) for the Office of Naval Research (ONR) during the period from April 1979 to March 1980. The purpose of this research is to extend the analytical taxonomies of Navy tactical command and control decisions developed by ADA during the previous phase of this research, develop a practical method for using taxonomies of decision characteristics and functions to evaluate and improve decision aids, and apply the method to the identification and (over)

DD FORM 1 JAN 73 1473 EDITION OF 1 NOV 65 IS OBSOLETE

Unclassified 393952

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

Block 20. (continued)

design of improved aids for antiair warfare (AAW) decisions. In addition, several analytical procedures have been developed to measure formally some of the characteristics and capabilities of decision aids that are outlined in the taxonomies.

TABLE OF CONTENTS

		Page
	Executive Summary	ES1
1	Introduction	1
	Part I: The Use of Taxonomies to Design and Evaluate Decision Aids	
2	A Procedure for Designing and Evaluating Decision Aids Based on a Set of Taxonomies	7
3	A Taxonomy of Information Processing and Decision Analytic Functions	17
4	The Characteristics of Antiair Warfare Decisions	25
5	Designing Improved Aids for Antiair Warfare Decisions: An Application of the Taxonomies	37
	Part II: Theoretical Extensions of the Taxonomies	
6	The Roles of Taxonomies and Analytical Procedures in Developing and Evaluating Decision Aids	67
7	A Review of Research on Limits to Human Information Processing and Decision Analytic Abilities	71
8	The Value and Credibility of Decision Aids	79
	References	101
	Distribution List	103

FIGURES

		Page
Figure 1	A Procedure For Evaluating Decision Aids	13
Figure 2	Displaying the Uncertainty in a Track's Position	52
Figure 3	Displaying the Uncertainty about the Existence of a Threat	53
Figure 4	Transforming a Lost Track into a Region where an Undetected Aircraft is Likely	55
Figure 5	Displaying Intelligence Estimates of the Direction and Type of Enemy Air Attack	57
Figure 6	Displaying the Likelihood that Defensive Weapons Can Destroy a Threat	59
Figure 7	Displaying the Likelihood that an Air Threat Can Penetrate to Any Location	60
Figure 8	Displaying a Summary of an Air Engagement Over Time	62
Figure 9	The Classical Paradigm	85
Figure 10	The Cuban Missile Crisis	88

TABLES

		Page
Table 1	Taxonomy of Decision Characteristics	8
Table 2	A Taxonomy of the Characteristics of Decision Aids	10
Table 3	A Taxonomy of Information Processing and Analytical Functions	18
Table 4	The Importance of Information Processing and Analytic Functions for AAW Decision Activities	39
Table 5	The Ability of Humans to Perform Information Processing and Analytic Functions for AAW Decision Activities	44

EXECUTIVE SUMMARY

This report describes research conducted by Applied Decision Analysis (ADA) for the Office of Naval Research (ONR) during the period from April 1979 to March 1980. The purpose of this research is to extend the analytical taxonomies of Navy tactical command and control decisions developed by ADA during the previous phase of this research, develop a practical method for using taxonomies of decision characteristics and functions to evaluate and improve decision aids, and apply the method to the identification and design of improved aids for antiair warfare (AAW) decisions. In addition, several analytical procedures have been developed to measure formally some of the characteristics and capabilities of decision aids that are outlined in the taxonomies.

A Procedure for Designing and Evaluating Decision Aids Based on a Set of Taxonomies

The approach developed in this research project uses several taxonomies to describe the characteristics of various command and control decisions, the characteristics of decision aids that could be used for command and control, and the information processing and analytic functions provided by people and decision aids to help a commander reach a decision. The taxonomy of information processing and analytic functions relies on the terminology of decision analysis and data processing. However, it is sufficiently general to describe intuitive and qualitative processes for dealing with a decision situation. The major information processing and analytical functions contained in the taxonomy are:

- acquire information,
- interpret information,
- store and recall information,
- structure and summarize information,

- analyze the decision situation,
- evaluate the decision, and
- transmit information and decisions.

By comparing the functions needed to analyze a particular decision with those that can be provided by a typical decision-maker and his staff, we can determine the functional requirements for an appropriate decision aid. These requirements are used to evaluate existing aids or guide the development of new aids.

The first step in using this procedure is to define carefully the decision situations for which an aid may be required. This includes identifying typical decision-makers and their level of command, the resources and options available to them, the organization that will help them analyze and implement decisions, the type of warfare and the nature of the adversaries they face, and the conditions under which the decisions must be made. Once the decision situations to be supported by an aid are defined, the characteristics of the decisions can be explored in detail. The characteristics provide a basis for assessing the functional abilities needed to analyze the decisions, and some of the characteristics (i.e., the frequency and importance of the decisions) are used to determine the value of aids that provide these functions.

The functional abilities of a typical decision-maker are compared with the requirements of a particular decision situation to produce a set of functional needs or deficiencies that a decision aid may be able to provide. The list of functional needs is a starting point for designing new aids, or a basis for evaluating existing aids.

The Characterisitics of Antiair Warfare Decisions

The procedure for designing and evaluating decision aids is applied to the range of decisions encountered in antiair warfare. The decisions considered here are primarily those encountered in conducting AAW operations at the unit level. They usually are made in a ship's combat information center (CIC) by a team of AAW personnel acting under the guidance of the ship's commanding officer and the task force commander. This guidance and the coordination provided by AAW personnel at the task force level constitutes another set of AAW decisions. These higher level decisions are discussed, but not considered in detail.

The decision-making activities that occur at the unit level are:

- detect and track aircraft and missiles;
- identify each track (i.e., determine the type of aircraft or missile, and whether it is friendly or hostile);
- assess the degree of danger posed by a threat (i.e., estimate its mission and the likelihood it will succeed);
- establish priorities for dealing with threats (i.e., which threats should be engaged first);
- assess the capabilities of alternative weapon systems for countering a threat (i.e., determine whether a weapon can intercept the threat and the likelihood that it will stop the threat);
- assign defensive weapons to counter each threat; and
- decide when, or under what conditions, to fire defensive weapons (e.g., fire a missile when the threat reaches a certain position).

In practice, these activities are treated as separate decisions, even though many of them are information processing activities that support the final decisions to assign a defensive weapon to an air target and fire the weapon. This is due to the fact that AAW operations are very complex and involve many individuals, each concerned with a portion of the overall problem. Each of the activities listed above could require a decision by different individuals during an air engagement. For instance, a tracker must decide whether or not to display a track based on possibly incomplete or conflicting radar signals, and an identification operator must decide what symbol to associate with the track (e.g., a symbol representing an enemy missile). These decisions are a method of assessing the information available to AAW personnel and communicating the assessments to others.

The taxonomies are used to explore the characteristics of each of these decisions. For instance, the outcomes relevant to all of the unit AAW decisions are dependent on a large number of uncertain factors. All AAW decisions are subject to review and approval, but, under the philosophy of command by negation, they are considered firm decisions until a higher level of command intervenes. The structure of current AAW decision activities at the unit level has been well defined by the Navy, and all are conducted using standard procedures. The quantity of information relevant to each of the decision activities processed in a ship's CIC during an engagement is extremely large. A variety of characteristics such as these are used to describe each of the AAW decisions in detail, and provide insight into the nature of decision aids that could support them.

Designing Improved Aids for Antiair Warfare Decisions: An Application of the Taxonomies

The AAW decision characteristics are used as a basis for identifying the importance of each information processing and analytic function for AAW decisions. The importance of these functions is then compared to how well AAW personnel can perform them without an automated decision-aiding system. The functional needs identified in this manner are:

- assess and communicate the uncertainty and credibility of information produced by the detection, tracking, assessment, and priority setting activities;

- combine the information available about threats and defensive weapons into an aggregate estimate of capabilities and intentions;
- store and recall subjective assessments of the status and capabilities of threats and defensive weapons;
- sort and categorize threats according to criteria appropriate to a specific combat situation;
- identify predefined patterns of data that indicate the existence and identity of tracks, and the capabilities of threats;
- identify trends in enemy tactics; and
- predict outcomes and evaluate alternative defensive actions.

The Navy currently uses an automated decision aid called the Naval Tactical Data System (NTDS) to provide some of these functions. The experience the Navy has gained from using NTDS provides a starting point and guide for attempts to design and implement improved aids for unit AAW decision-making activities.

NTDS is a specialized information processing system that accepts real-time data from radars, aircraft, weapons systems, operators, and other ships; processes this information and displays it on small-screen consoles; exchanges the information via digital data links with other NTDS units (or units with compatible data systems); and reports some of the information to other commands and non-NTDS units via teletype. NTDS automates many of the routine information processing activities conducted in manual AAW operations (e.g., plotting tracks, computing intercepts, communicating estimates and decisions) with the result that CIC personnel can perform these tasks more rapidly and for a larger number of tracks. NTDS is very good at accepting, storing, and displaying data in certain standard formats, but there is little flexibility to process information that cannot be expressed using these formats.

The primary control and output device for this sytem is an NTDS console, which consists of a small screen resembling a radar repeater and a variety of control keys. The information displayed on an NTDS console is a set of "synthetic" symbols generated by the NTDS computers. Although radar signals can also be displayed directly on NTDS consoles, this raw data is not transmitted among NTDS units or used to allocate weapons. Information associated with NTDS symbols can be retrieved by console operators and displayed on panels near the screen. Digital communications through the NTDS computers is augmented by voice channels, both within a CIC and among ships and aircraft.

Translating radar signals into information that can be processed by NTDS requires a sequence of subjective judgements that are neither supported nor adequately documented by NTDS. The process is very tedious and is performed by enlisted men who are trained on the job. However, the quality of all subsequent decisions depends on their judgements. NTDS provides no way for them to communicate the confidence they place in their judgements, and no way for others to review at a later time the raw data from which the judgements were made. Inaccurate data and incorrect interpretations of incomplete data are rapidly disseminated to other NTDS users and the source of an error often is difficult to identify.

The process of entering information and instructions into an NTDS console is slow relative to verbal communications. Communicating rapidly via a set of function keys is not a natural activity for most individuals, especially in a crisis. As a result, information often reaches CIC personnel via voice circuits before it shows up on an NTDS console.

Probably the most fundamental improvement that could be made in AAW decision aids is to provide them with the capability of accepting, processing, and displaying uncertainty. Many of the other functional improvements depend on implementation of this capability. The difficulty lies in defining simple ways to enter and display uncertainties that will not require too much time of AAW personnel or overwhelm them with too much information.

Would the inclusion of uncertainty in AAW data overwhelm decision-makers with too much information? There is already a potentially overwhelming amount of data in NTDS, and yet AAW personnel still find it necessary to consider and mentally process uncertainties in order to reach a decision. If some of this processing were done by a decision aid and displayed in a simple format, the decision-maker could devote his attention to interpreting the results rather than trying to calculate them.

An aid capable of processing uncertain data could play a more significant role in detecting, tracking, and identifying aircraft and missiles. By matching data received about tracks with the known characteristics of various aircraft and missiles, an automated aid could calculate the likelihood that the data is caused by various offensive and defensive weapons systems, even if some of the data is missing or contradictory. The same type of information processing would indicate the probability that data was generated by <u>any</u> aircraft or missile, thus helping the user detect tracks from noisy data.

Once AAW decision aids have the capability of processing subjective estimates of uncertainty, they can be given the flexibility to accept and use prior estimates of enemy tactics and intentions from intelligence sources or CIC personnel.

Another major area where AAW decision aids could be improved is the generation of better summaries of the tactical situation. NTDS does very little to aggregate or summarize the detailed information it contains, forcing users to do the necessary information processing. An automated aid that summarizes

AAW information should display the result graphically to help user interpret it quickly. Several different displays are described that would be appropriate for summarizing different aspects of AAW operations.

One area where improvements have been attempted in existing aids is the evaluation of alternative allocations of defensive weapons. These attempts use simple algorithms to recommend or make weapon assignments, but do not allow the user much interaction with the evaluation process. There are many possibilities for improved aids to evaluate AAW resource allocations. However, their design depends on other improvements in the information processing and analytic abilities of AAW decision aids. Attempts to design better resource allocation aids are likely to meet the same fate as current systems if they are not preceded by improvements in an aid's ability to interact with the decision-making process, and capture in a simple but realistic way the uncertainties and subjective information relevant to his decision.

The Roles of Taxonomies and Analytical Procedures in Developing and Evaluating Decision Aids

In the course of our research we found the need for more precise analytical procedures than taxonomies to evaluate at least some of the characteristics of decision aids. Taxonomies provide a useful and flexible framework for conducting a general evaluation of decision aids. However, they should be augmented by more precise analytical procedures to guide specific resource allocations.

Taxonomies represent a general conceptual approach for thinking about decision aids. They have several important advantages that guarantee that they will remain useful even after more precise evaluation procedures have been developed. Taxonomies provide a comprehensive checklist to ensure that all relevant aspects of an aid and the decisions it supports are considered. They can also organize the evaluation process into a series of logical steps by showing the sequence in which issues should be considered and compared. Taxonomies also provide a framework for decomposing general characteristics of decision situations and aids into basic elements that can be assessed or measured. This decomposition provides a guide for establishing the relative importance of desired characteristics. Finally, one of the major strengths of taxonomies is that they do not require precise, unambiguous assessments of the characteristics they contain, thus making it easy to use them.

However, the ambiguity permitted by a taxonomy is also one of its major weaknesses. It may not be possible to interpret correctly imprecise assessments of an aid's characteristics or capabilities, or to determine the implications of these assessments for the applicability of the aid in various decision situations. Taxonomies can provide a guide to evaluating aids, but not an operational measure of their value.

We anticipate that taxonomies and analytical procedures eventually will complement each other, and will be used together to evaluate decision aids. An evaluation procedure using both would probably start with taxonomies to make sure the important issues or features associated with an aid are considered and to identify the key tradeoffs that must be made. Then one or more rigorous analytical procedures would be used to produce formal measures of an aid's potential value. Finally, general evaluation procedures based on taxonomies would be employed to deal with aspects of an aid or decision situation that are poorly defined or for which no formal evaluation procedure exists.

The Value and Credibility of Decision Aids

In the traditional view, the decision-maker must decide to accept or reject the conclusions produced by an aid. If he rejects the aid's results, he makes a decision based on intuition; if he accepts the results, he chooses the aid's conclusions. There are at least two problems with this. First, there are no explicit criteria for determining whether to accept or reject the aid's results. The decision-maker is often placed in the position of having to judge the results of an aid whose technical details he doesn't really understand. Second, and perhaps more important, in rejecting either the aid or his intuition, the decision-maker is throwing away potentially useful information. The danger of neglecting intuition is the danger of neglecting important insights just because they are hard to model; the danger of neglecting the aid is the danger of neglecting valid information and logic just because it is hard to understand.

A broader and more useful view of a decision aid is that it provides a decision-maker with additional data on which to base his choices. The result of decision-aiding is information for the decision-maker, not a substitute for intuition, good judgement, or clear thinking. Viewing analytical aids as complements to, rather than substitutes for, unaided intuition frees us from the unnecessarily rigid vision of formal analysis as a competitor of the intuitive process.

In effect we are claiming that an aid should be viewed as a special kind of expert. No one would ever argue that all expert opinions should be believed, or that all expert advice should be followed to the letter. It makes just as little sense to treat the conclusions derived from an aid as inviolate. However, it is equally shortsighted to completely ignore an aid if it can provide insights, but fails to provide a perfectly credible answer. The information produced by the aid, just like the advice of the expert, must be tempered by the decision-maker's feelings about the credibility of the analysis that produced it.

An aid's credibility is related to the decision-maker's feelings about how future modeling and information could affect the aid's results. Furthermore, the stability of the aid's results is a function of the stability of its inputs and assumptions. If all the inputs and assumptions were invariant to future information or modeling, then the aid would be perfectly credible. The credibility of an aid can be defined formally in terms of the results it would produce if it were developed to the point where further modeling or information would not change it or its results. This definition is the basis of a formal analytical procedure for determining the credibility of a decision aid. The steps required by the procedure are specified for aids based on decision trees or models with a finite number of inputs.

1. INTRODUCTION

This report describes research conducted by Applied Decision Analysis (ADA) for the Office of Naval Research (ONR) during the period from April 1979 to March 1980. The purpose of this research is to extend the analytical taxonomies of Navy tactical command and control decisions developed by ADA during the previous phase of this research, develop a practical method for using taxonomies of decision characteristics and functions to evaluate and improve decision aids, and apply the method to the identification and design of improved aids for antiair warfare (AAW) decisions at the unit and task-force levels. In addition, several analytical procedures have been developed to measure formally some of the characteristics and capabilities of decision aids that are outlined in the taxonomies. Some of these, including a method for measuring the credibility of decision aids, are described in this report.

The approach developed in this research project uses several taxonomies to describe the characteristics of various command and control decisions, the characteristics of decision aids that could be used for command and control, and the information processing and analytic functions provided by people and decision aids to help a commander reach a decision. By comparing the functions needed to analyze a particular decision with those that can be provided by a typical decision-maker and his staff, we can determine the functional requirements for an appropriate decision aid. These requirements are used to evaluate existing aids or guide the development of new aids.

This approach is applied to the range of decisions encountered in unit level antiair warfare, and several functional areas are identified where AAW decisions could benefit from improved aids. The capabilities of existing AAW decision aids to provide the needed functions are discussed, and ways in which functional improvements could be designed and implemented are explored.

Decision aids are viewed as sources of information that complement, rather than replace, intuition and judgement. In this report decision aids are defined broadly to include simple information processing devices as well as sophisticated systems for interpreting information and projecting the consequences of alternative courses of action.

The decision aid evaluations and designs described here use the taxonomies of decision and decision aid characteristics developed during the previous phase of this research. These taxonomies are described in detail in (7), and summarized in Chapter 2 of this report. Chapter 3 proposes an additional taxonomy of information processing and analytic functions that are common to both decision-makers and decision aids. Chapter 4 describes the range of decisions associated with AAW operations using the taxonomy of decision charateristics. Chapter 5 applies the taxonomies to the information processing functions needed for AAW decisions, and explores the design of improved aids for AAW operations. Chapter 6 consideres the limitations of taxonomies for evaluating and designing decision aids, and discusses the role of more precise analytical procedures. Chapter 7 reviews previous research to analyze the limits of human decision making activities and the role of decision aids in overcoming these limits. Chapter 8 describes several quantitative analytical procedures for measuring the characteristics of decision aids, with emphasis on determining and improving their credibility.

The report is divided into two parts. Part I describes the taxonomies and shows how they can be applied to AAW decisions. This part of the report is self contained. Readers primarily interested in improving AAW decision aids or applying taxonomies to other decision situations can limit their attention to Part I.

Part II deals with the theoretical basis for elements of the taxonomies. It describes formal, quantitative methods for measuring the value and credibility of decision aids. These methods are designed to complement and extend the taxonomies by providing precise, umambigious ways to deal with important aspects of the design and evaluation process.

PART I

THE USE OF TAXONOMIES TO DESIGN AND EVALUATE DECISION AIDS

2. A PROCEDURE FOR DESIGNING AND EVALUATING DECISION AIDS BASED ON A SET OF TAXONOMIES

Different taxonomies are needed to deal with various aspects of decision making and aiding. In particular, different taxonomies are appropriate for describing the features of decision situations, decision-makers, and decision aids. Taxonomies have also been developed for the types of decision situations encountered in tactical command and control, and the analytical methods appropriate for different classes of decisions. See (4) and (7).

The previous phase of this research produced three closely related taxonomies. The first taxonomy groups Navy tactical command and control decisions into various categories. Representative decisions can be specified for each category, although the categories are sufficiently broad to include dissimilar decision situations. As a result, this taxonomy is useful mainly for identifying the range of decisions an aid may have to support. The second taxonomy contains a set of attributes or analytic measures that can be used to characterize a decision situation. These attributes are specified in both technical and intuitive terms, and a scale is defined for assessing each attribute. These attributes include: the number of decision strategies available to the decision-maker, the level of resources involved, the amount of information available and the time period within which the decision must be made. Table 1 contains a summary of the taxonomy of decision characteristics. The third taxonomy, which is not as fully developed as the others, specifies the characteristics of decision-aiding systems in terms that can be compared with the corresponding characteristics of the decisions for which they will be used. This taxonomy deals primarily with an aid's physical characteristics and capabilities, such as its cost, support requirements, and data processing and storage capabilities. Table 2 contains a representative taxonomy of decision aid characteristics.

TABLE 1: TAXONOMY OF DECISION CHARACTERISTICS

1. The Decision-maker's Resources

   Equivalent force level controlled by the decision-maker

2. The Importance of Decision to Decision-maker

   Importance of the decision relative to the other decisions made by the decision-maker over one year.

3. The Number of Decision Strategies

   Number of alternatives available in the primary decision problem (i.e., first decision node in a decision tree)

4. The Number of Significant Factors

   Number of factors (i.e., state variables) that could have a significant impact on the outcome

5. The Number of Outcome Attributes

   Minimum number of outcome variables (i.e., attributes that must be considered to adequately represent outcomes

6. Outcome Measurability

   Percentage of outcome variables (attributes) requiring a subjective scale

7. Contingent Decisions

   The importance of contingent planning in this situation

8. Probabilistic Dependence

   Average or typical number of variables that have a significant and direct impact on each outcome variable

TABLE 1: TAXONOMY OF DECISION CHARACTERISTICS (Cont'd.)

9. The Degree of Risk

   Importance of low-probability, high-consequence events

10. Review and Approval

    Extent of required review and approval

11. Structural Uniqueness

    Extent to which existing plans or procedures can be used to deal with the decision

12. The Quantity of Information

    Number of messages related to decision received per day

13. The Variability of Information Value

    Percentage of messages that are significantly more valuable than the average

14. The Reliability of Information Sources

    Percentage of information sources considered reliable

15. The Time Available for the Decision

    Time from recognition of a decision to the point where an action must be taken

16. The Frequency of Decision

    Mean time between recurrence of the decision

TABLE 2: A REPRESENTATIVE TAXONOMY OF THE CHARACTERISTICS OF DECISION AIDS

1. Unit cost of the aid, including a proportional share of the development costs.

2. Cost of using the aid, including the level of effort required to use or program it.

3. Support requirements, including data sources, other aids and equipment, and physical space.

4. Reliability, including redundancy, self-monitoring, and a capability for graceful degradation.

5. Data processing capability, as measured by the number of calculations and amount of data that can be processed per unit time.

6. Data storage capability, including both rapid-access and slow-access storage.

7. Capability to maintain data security, including access control and encoding.

8. Data verification capability, including error checking and cross checking data from multiple sources.

9. Communications capability, including data transmission rates and number of communication channels.

10. Ability to prioritize its own operations, including the ability to interrupt and restart a procedure.

11. Facilities for testing and updating algorithms, as measured by the ease with which procedures or analyses can be restructured.

12. Ability to monitor and update information, including information sorting and screening.

13. Complexity of the aid-user interface, including the level of training required of the user, the sophistication of the algorithm, and the extent to which the aid summarizes and supplies its outputs.

14. Compatibility with existing systems and procedures.

The taxonomies of decision characteristics and decision-aid characteristics are closely related. Each decision characteristic has implications for the features of a decision aid that would be appropriate for that decision. A qualitative mapping between the two taxonomies was developed and demonstrated during the previous phase of this research. For example, the decision characteristic measuring the time available for a decision has an effect on four of the aid characteristics: cost of use, reliability, data processing capability, and complexity of the user interface. Conversely, the decision aid characteristic measuring the cost of use is related to three decision characteristics: importance of the decision to the decision-maker, structural uniqueness, and the time available for a decision.

However, these linkages between the characteristics of decisions and decision aids are not well defined, and the process of designing or evaluating an aid using the taxonomies in Tables 1 and 2 requires a considerable amount of judgement on the part of the user. To help overcome this problem, we have developed an additional taxonomy that describes the information processing and analytical functions that are provided by a decision aid and needed to examine a decision situation. Because this taxonomy specifies a common set of functional requirements for decision aids and decision situations, it can be used to relate the characteristics of decision aids to those of the decisions they support. It also can be used to describe the analytic functions provided by a decision-maker and his staff, which means that aids can be designed to perform functions that do not duplicate those already available.

Although the taxonomy of analytic functions relies on the terminology of decision analysis and data processing, it is sufficiently general to describe intuitive and qualitative processes for dealing with a decision situation. The taxonomy contains information processing and analytical functions such as:

sorting and reorganizing data, identifying patterns and relationships, using an analysis to predict outcomes, and identifying the most significant elements of a decision. This taxonomy is discussed in detail in the next chapter.

The Procedure

The manner in which the taxonomies can be used to design or evaluate a decision aid is outlined in Figure 1. The sequence of assessments and logic is shown as a flow chart, starting with either a specific decision aid or a set of decisions for which an aid is needed, and proceeding through a series of steps to an assessment of the functions the aid should provide. The required functions can be used as a guide for designing an appropriate aid, and as a basis for assessing the value of an aid. The steps shown in Figure 1 make use of several taxonomies to describe and compare the features and capabilities of various aspects of the decision process. Some of the steps are optional; they are shown with dashed lines. These steps can be bypassed if the characteristics of the decision aid and its potential users are well understood, but doing so will make the evaluation and design process less explicit and more dependent on the intuition of the analyst. If all of the steps in Figure 1 are carried out, the process makes use of four taxonomies: three that describe the characteristics of decision aids, decision-makers, and decision situations; and one that specifies the information processing and analytic functions used in decision making. The last taxonomy is used to compare the requirements of decision situations with the capabilities of decision-makers and aids.

The first step in using the procedure outlined in Figure 1 is to define carefully the decision situations for which an aid may be required. This includes identifying typical decision-makers and their level of command, the resources and options available to them, the organization that will help them analyze and implement decisions, the type of warfare and the nature of the adversaries they face, and the conditions under which the decisions must be

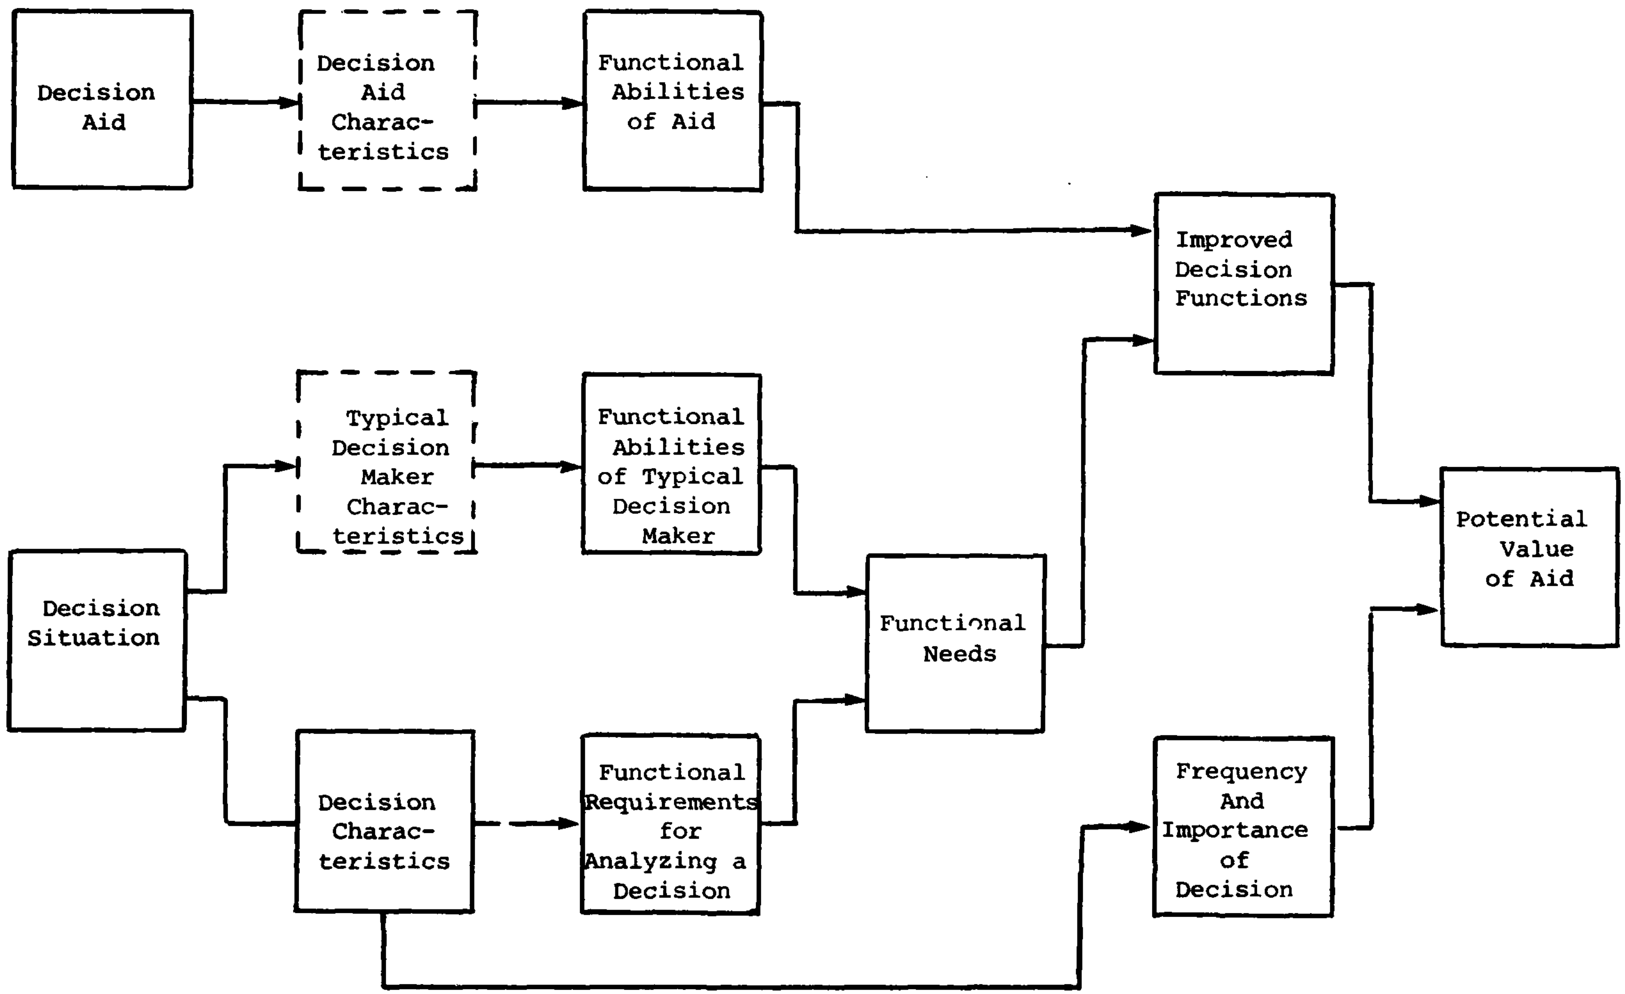

A PROCEDURE FOR EVALUATING DECISION AIDS

Figure 1

made. A significant part of defining the decision is bounding the range of problems to be considered. Attempts to design or evaluate aids for a very broad class of decisions are likely to result in large, expensive, and overly complex aids, or aids that are evaluated in terms of the wrong set of decisions.

Once the decision situations to be supported by an aid are defined and bounded, the characteristics of the decisions can be explored in detail. The taxonomy shown in Table 1 can be used to accomplish this task. The taxonomy acts as a checklist to ensure that all of the decision characteristics relevant to the design and evaluation of decision aids are considered. The characteristics provide a basis for assessing the functional abilities needed to analyze the decisions, and some of the characteristics (i.e., the frequency and importance of the decisions) are used to determine the value of aids that provide these functions.

A similar taxonomy could be used to specify the characteristics of decision-makers who might use the aid. A taxonomy of decision-maker characteristics has not been developed, but it is probably not necessary to use one to describe the officers who make Navy tactical command and control decisions. Part of the problem with trying to specify the characteristics of decision-makers in much detail is that Navy officers differ in such characteristics as their level of training in analytical methods and prior experience with decision aids. However, other characteristics, such as the organizational structure within which the decision-maker must operate, are essentially the same for each Navy officer responsible for the same class of decisions. For the purpose of evaluating decision aids for tactical command and control, one should assume that they will be used in a manner specified by Navy procedures and that the individuals using them will have a level of training and skill typical of Navy officers.

Once the characteristics of decision-makers and decision situations have been assessed, either directly or with the help of taxonomies, they are translated into a common set of functional abilities and requirements. This translation process is done intuitively, by considering the implications of each characteristic for a particular function. The assessment of functional abilities or requirements is done separately for decision-makers and decision situations, and then compared to identify functional needs. Since the decision-maker is specified only as being a typical Navy officer, his functional abilities can be assessed once for each level of command. Only functional requirements need to be reassessed when new decision situations are considered.

After functional abilities and requirements are assessed, the comparison process is relatively straightforward. Comparing the functional abilities of a typical decision-maker with the requirements of a particular decision situation produces a set of functional needs or deficiencies that a decision aid may be able to provide. The list of functional needs is a starting point for designing new aids, or a basis for evaluating exis.ing aids.

If an aid has been designed, its characteristics should be specified in as much detail as possible. The taxonomy shown in Table 2 can be used to accomplish this task. The aid's characteristics are used as a guide for assessing its functional abilities, expressed in the same terms as those used for decision-makers and decision aids. Comparing an aid's functional abilities with the functions needed to analyze the decisions it supports indicates the extent to which the aid can improve the decision process. The improvements in information processing and analysis are evaluated in terms of the frequency and importance of the decision situation. The resulting estimate of an aid's potential value would be realized if the aid is used whenever the decision situation under consideration is encountered. It may be necessary to discount this potential value if it is believed that not all decision-makers are willing to use the aid, or if the aid is not available every time a relevant situation arises.

3. A TAXONOMY OF INFORMATION PROCESSING AND DECISION ANALYTIC FUNCTIONS

This chapter describes a taxonomy of information processing and decision analytic functions. The taxonomy is used to relate the characteristics of decision aids to those of decision situations and decision-makers. It is intended to be sufficiently general to describe the abilities and requirements of aids, decision-makers, and decisions in common terms.

The functional taxonomy is based on the assumption that some attempt is made by the decision-maker to use relevant information to understand the nature and consequences of a decision. In other words, it is assumed that some information processing or analysis (perhaps at an intuitive level) is needed to make the decision. If this assumption is not valid, then there is little point in developing a decision aid.

Although rather perfunctory information processing tasks are included in the taxonomy, they are part of the process of reaching a decision. The taxomomy is not an attempt to detail every aspect of information processing, but only those functions that are directly relevant to decision making. Also, the taxonomy does not specify how the functions are carried out, but only what they accomplish.

The taxonomy is outlined in Table 3. It starts with the task of acquiring information and proceeds through a series of more complex functions until an analysis of the decision is integrated with existing knowledge, a choice is made, and the conclusions drawn from the analysis are transmitted to other individuals or organizations. The order in which the information processing and analytic functions are specified in the taxonomy often does not correspond to the order in which these functions are performed. In fact, a typical decision process can

TABLE 3: A TAXONOMY OF INFORMATION PROCESSING AND ANALYTIC FUNCTIONS

1. Acquire information

   A. Observe and measure physical phenomena (e.g., radar signals)

   B. Monitor information channels (e.g., communications from another command)

2. Interpret information

   A. Assess uncertain quantities (e.g., estimate of enemy intentions)

   B. Estimate aggregate descriptors (e.g., a measure of combat readiness)

3. Store and recall information (i.e., no change or reorganization of data)

4. Structure and summarize information

   A. Aggregate or collect information into pre-established categories (e.g., important vs. routine)

   B. Sort or arrange data in a specified order and format (e.g., rank air threats in terms of proximity to task force)

   C. Compare and filter data (i.e., eliminate portions of information)

   D. Recognize pre-established patterns of information (e.g., look for flight characteristics of cruise missile using a prestructed model or template)

5. Analyze decision situation (i.e., decompose decision into basic elements)

   A. Identify and define significant elements of situation; search for new elements

      (1) Alternatives
      (2) Factors and issues affecting the decision
      (3) Uncertainties and information sources
      (4) Outcomes and preferences

TABLE 3: A TAXONOMY OF INFORMATION PROCESSING AND ANALYTIC FUNCTIONS (Cont'd.)

B. Identify patterns and relationships

(1) Recognize trends or patterns in data
(2) Define causal links (i.e., dependencies) among elements of situation
(3) Fit functional relationships to data (correlate data); update relationships based on new data

C. Search for alternative representations of decision situation (i.e., Is there an easier or more insightful way to look at th problem?)

6. Evaluate decision situation

A. Determine implications of analysis

(1) Calculate summary data (e.g., compute total number of aircraft available)
(2) Use logic from an analysis to predict outcomes and extrapolate data
(3) Determine optimum alternatives
(4) Calculate sensitivity of outcomes and decisions to changes in data or assumptions
(5) Determine value of additional information and cost of delay

B. Integrate the results of an analysis with existing knowledge and intuition

(1) Identify the most important elements and relationships
(2) Develop a simple explanation of analytic results (that can be internalized by decision-maker); interpret and consolidate results of analysis
(3) Update decision-maker's internal model of decision situation

C. Select an alternative

7. Transmit information and decisions (i.e., instructions)

A. Verbal

B. Alpha-numeric/digital

C. Graphical

skip back and forth among several of the functions. The sequence of tasks contained in the taxonomy is intended only to suggest the general flow of analytical activity as a decision is reached.

The first function is acquiring information. The nature of this task depends on whether the source of the information is a data processing or storage system (e.g., another command center, a computer, or a manual) or a physical measurement. If the information has been preprocessed and stored or transmitted by another individual or system, the process of acquiring the information requires only that is be successfully received and decoded. However, direct acquisition of physical data requires some sort of measurement procedure, usually based on sensors and displays.

The second function is interpreting information. This function requires a judgement on the part of the decision-maker or his staff to assess the uncertainty inherent in a piece of information, or interpret its meaning in terms of some aggregate parameter (e.g. combat readiness). This type of information processing is done without recourse to a formal or explict analysis of the data, although the estimates may be based on an intuitive model of the situation. The process of interpreting the results of an explicit calculation or analysis is considered separately in another part of the taxonomy. The function of interpreting information relies on methods for expressing, testing, and comprehending subjective estimates (e.g. scales for expressing degree of belief in a statement, or debiasing procedures for subjective probability assessments).

The third function is the storage and retrieval of information. This function does not involve any change or reorganization of the data, other than placing the information in a form that can be stored or interpreted after retrieval. For example, a computer may translate a message into a series of bits that can be stored on a disk, but when the message is recalled it will be essentially the same as when it was stored. Any reorganization of the data during the storage and retrieval process is considered part of the next function.

The fourth function, structuring and summarizing information, includes the many ways that information can be organized and placed in a form that is meaningful to the decision-maker. These procedures generally involve comparing the data to a pre-established pattern or criterion, or using the data to recognize when one exists. For instance, pieces of information can be grouped into a set of categories based on a rule for discriminating between categories (e.g., incoming messages can be classified as urgent or routine). Alternatively, information can be reorganized by sorting it according to a particular sequence or format (e.g., ranking air threats according to their proximity to a task force). Both of these operations change the order in which data is presented to the decision-maker without changing the total amount of information. In contrast, information can be filtered to eliminate data that does not meet established criteria. A more sophisticated method of classifying information occurs when data is scanned to determine whether or not a designated pattern exists (e.g., examing the flight characteristics of an air threat to see if it is a cruise missile). The criteria used to reorganize data is fixed in each of these procedures. Developing a new criterion (i.e., a new model of the relationships among data elements) is part of the next function.

The fifth function is analyzing a decision situation. The process of recognizing and specifying the basic components of a decision is probably the most difficult and creative part of the decision process. Unfortunately, there are relatively few decision aids to help someone carry out this task. The outcome of a decision situation often hinges on how well the individuals involved have identified the options available to them and the factors that have a significant bearing on the decision. Some of these are likely to be uncertain, so it may be necessary to identify various options for gathering information to reduce the uncertainty. Any decision-making activity requires some consideration (either implicit or explicit) of the outcomes that could result from alternative courses of action, and the relative value of each outcome.

In addition to recognizing the elements of the decision problem, it is usually necessary to identify the relationships among them. This typically is done by recognizing trends or patterns in the data, and assuming that these patterns can be explained by dependencies among the elements of the decision. These causal links are often the most important and least recognized assumptions in an analysis. They can be established by fitting parametric relationships to the data, and updating the relationships as new data is received. A powerful way to analyze a decision situation is to search for alternative ways to think about the problem. Comparing different models of the same situation often provides insight about the most significant elements of the problem.

The sixth function is evaluating a decision situation using the results of an analysis, and combining the results with existing knowledge and intuition. Once a decision situation has been decomposed into a set of elements and relationships, logic can be used to determine the implications of the analysis. The simplest form of logic is the calculation of some summary information, such as the total amount of fuel in storage. Simple calculations like this require very little modeling beyond an assumption that the quantity being calculated is relevant to the decision. Usually more sophisticated logic is required to determine the implications of an analysis. Often it includes some sort of extrapolation procedure to predict the outcomes resulting from a decision, and it may include calculations of the uncertainty associated with these outcomes. In addition, the logic used to predict outcomes can be imbedded in an optimization procedure to determine the course of action that produces the most desirable outcomes. Since the recommended course of action depends on all the assumptions contained in the analysis, additional calculations often are used to test the sensitivity of the predicted outcomes and recommended decisions to changes in data and assumptions. If the analysis is probabilistic or dynamic, further logic can be used to determine the value of additional information or the cost of delay.

Numerous computer programs have been written to calculate the implications of an analysis. However, few are sufficiently general to be useful for more than a few specific decisions. General purpose algorithms have been developed for special classes of problems (e.g., linear programming), but they often must be adapted or tailored to fit a specific decision. In addition, they usually deal with only part of the problem, so they become part of a larger analysis.

After an analysis has been completed, it must be integrated with existing knowledge and intuition. This function is often done intuitively with little support, although it is one of the most important. Analyses are rarely accepted without reconciling their results with the decision-maker's current understanding of the situation. The extent to which a decision-maker's current view of the problem is updated or replaced depends on the credibility and limitations of the analysis. (A detailed discussion of the concept of credibility is contained in Chapter 8). Often the purpose of an analysis is to give the decision-maker a few basic insights that then become the basis for a decision. This can be accomplished by identifying the most important elements and relationships in an analysis, and using them to form a relatively simple model of the decision situation that can be internalized by the decision-maker.

The last function is transmitting the results of the decision process to other individuals and organizations. These results can be instructions or decisions, or simply information that results from data processing and analysis. The form of the communications can be verbal, alpha-numeric, or graphical depending on the nature of information being transmitted. Typically a decision-maker has a variety of aids to support this function.

4. THE CHARACTERISTICS OF ANTIAIR WARFARE DECISIONS

In this chapter and the following one, we apply the procedure for designing decision aids that was described in the previous chapters to the command and control decisions of antiair warfare (AAW). The decisions considered here are primarily those encountered in conducting AAW operations at the unit level. They usually are made in a ship's combat information center (CIC) by a team of AAW personnel acting under the guidance of the ship's commanding officer and the task force commander. This guidance and the coordination provided by AAW personnel at the task force level constitutes another set of AAW decisions. These higher level decisions are discussed, but not considered in detail.

This chapter describes the decisions associated with AAW operations and uses the taxonomy in Table 1 to explore their characteristics. The next chapter uses this information to explore the relative importance to the AAW decision process of each decision function described in the previous chapter, and the nature of the decision aids that could help AAW personnel perform the more important functions.

The AAW Decision Environment

The previous phase of this research developed a taxonomy for categorizing Navy command and control decisions along four dimensions: level of command, type of warfare, decision function, and decision context. In terms of this classification, the decisions under consideration here are made at the unit level and, to a lesser extent, by the task force commander and his staff. They deal entirely with air defense. (In particular, the capabilities of NTDS to support decisions for surface and antisubmarine warfare are not considered here.) They require the decision-makers to position and allocate assets to accomplish specific AAW tasks. Most of the decisions are made in the context of executing an air defense plan specified by higher levels of command, although some decisions at the task force level may be incorporated in that plan.

By restricting our attention to this rather small subset of the decisions encountered in Navy command and control, it is feasible to describe each of the major decision-making activities involved. These decision-making activities at the task force level deal primarily with assigning AAW responsibilities and assets to individual units, and then coordinating their activities. A task force commander can delegate some of his decisions to a sector AAW coordinator who is responsible for air defense in a designated sector. The sector AAW coordinator functions in the same way that the task force commander's staff would within his sector, so his decision activities will be considered as part of those at the task force level.

The major decision-making activities that occur at the task force level are:

- establish operational guidelines for AAW personnel and interpret rules of engagement specified by higher levels of command (e.g., define the conditions under which a unit can fire on an air target without further authorization);
- interpret intelligence reports and anticipate the type of air attack that might occur;
- position ships and aircraft to maximize the chances of detecting and intercepting air threats;
- establish controls on electronic emissions (EMCON) and select electronic warfare tactics;
- designate air defense sectors and assign responsibility for AAW operations within a sector to a ship or a sector AAW coordinator;
- assign defensive aircraft (interceptors and electronic surveillance) to sectors;

- establish procedures for the movement of friendly aircraft near the task force; and

- resolve conflicting decisions or actions by AAW units and coordinate their activities.

The last decision-making activity in this list is made both before and during an engagement by one of several officers designated to coordinate air defense activities. The principal officer in this category is the antiair warfare coordinator (AAWC). He is supported by a force track coordinator (FTC) who monitors and reconciles the tracking and identification decisions made by units in each sector. Since the decisions made by these officers are closely related to those made at the unit level, they currently are supported by the same decision aids. However, these aids provide little direct support for the other AAW decision activities in the list above.

The task force commander and his staff are responsible for the first five decision-making activities listed above, and for coordinating the activities of the AAWC with those of officers concerned with other types of warfare, such as the surface action commander, the antisubmarine warfare commander, and the electronic warfare coordinator. The AAWC and FTC are located in a combat information center (CIC) where they can monitor the activities of individual AAW units. These officers are primarily responsible for the last three decision activities listed above. They report to the task force commander and his immediate staff, who are located in their own command center, which rarely contains automated decision aids.

The decision-making activities that occur at the unit level are supported by fairly sophisticated aids. These activities are:

- detect and track aircraft and missiles;

- identify each track (i.e., determine the type of aircraft or missile, and whether it is friendly or hostile);

- assess the degree of danger posed by a threat (i.e., estimate its mission and the likelihood it will succeed);

- establish priorities for dealing with threats (i.e., which threats should be engaged first);

- assess the capabilities of alternative weapon systems for countering a threat (i.e., determine whether a weapon can intercept the threat and the likelihood that it will stop the threat);

- assign defensive weapons to counter each threat; and

- decide when, or under what conditions, to fire defensive weapons (e.g., fire a missile when the threat reaches a certain position).

In practice, these activities are treated as separate decisions, even though many of them are information processing activities that support the final decisions to assign a defensive weapon to an air target and fire the weapon. This is due to the fact that AAW operations are very complex and involve many individuals, each concerned with a portion of the overall problem. Each of the activities listed above could require a decision by different individuals during an air engagement. For instance, a tracker must decide whether or not to display a track based on possibly incomplete or conflicting radar signals, and an identification operator must decide what symbol to associate with the track (e.g., a symbol representing an enemy missile). These decisions are a method of assessing the information available to AAW personnel and communicating the assessments to others.

Since there are significant uncertainties associated with the information received by AAW personnel, it is possible for them to make different assessments and decisions based on the same data. When this happens, the decision is referred to a higher level of command (e.g., the force track coordinator) where the final decision is made. Unfortunately, this process of treating information processing and assessment activities as separate decisions tends to suppress some of the most important information needed to assign and fire a weapon: the uncertainty and credibility associated with the data produced by the earlier decisions. For instance, "deciding" that a threat is a particular type of enemy aircraft can lead to some very inappropriate actions if the threat is actually something else. In spite of this difficulty, we will treat each of the activities listed above as a separate decision in the discussion that follows and then return to the problem of improving the flow of information among AAW decision-makers.

At the unit level, all of the decision-making activities take place in a CIC, where they are supported by either automated or manual decision aids. The commanding officer (CO) of the ship has the responsibility for defending it from all threats, including aircraft and missiles. However, since he also has many other responsibilities, he can delegate AAW decisions to personnel located in the CIC. His principal representative there is the tactical action officer (TAO). The TAO generally has authority to fire weapons without seeking approval, when necessary for the immediate defense of the ship. During combat he is trained to manage all unit AAW operations and coordinate his actions directly with the AAWC.

The remaining AAW personnel in a ship's CIC report to the TAO. Their responsibilities correspond to the unit level decision activities listed above. Enlisted men are assigned jobs as detector/trackers, and usually are trained on the job in a few days. On a ship with an automated AAW system, the job consists of watching radar screens and assigning synthetic (i.e., computer generated) symbols to the blips caused by radar returns. As the radar signals move across the screen, the detector/trackers must instruct the computer to move the synthetic symbols along the same paths. ID operators use a variety of information to determine the identity of each object on the screen, and modify the synthetic symbols so others can see what they represent. ID operators can also enter alphanumeric data for each track that can be retrieved by other CIC personnel.

The detector/trackers and ID operators work under the direction of a unit track supervisor, who usually has more experience than they do. The track supervisor assigns responsibility for various tracks to individual trackers, resolves conflicting estimates or actions on their part, and works directly with the force track coordinator. The detector/trackers and ID operators distill the flow of information coming into the CIC and decide how it should be summarized and stored for use by other CIC personnel.

Assessing and prioritizing the threats represented by the synthetic symbols is primarily the job of the TAO and, if the ship has an automated AAW system, the ship's weapons coordinator (SWC). The SWC sits at a computer console while the TAO generally stands behind him and consults with him to determine the best way to deal with air threats. The SWC also keeps track of the availability of defensive weapons systems and helps the TAO assess their capability to destroy air threats. Defensive weapons can be assigned to threats by either the SWC or the TAO, with the TAO exercising veto power if he does not accept the SWC's decisions. Working with the SWC may be several specialized weapon controllers, including an intercept controller, fire control system controller, and engagement controller, depending on the size of the ship and the number of stations in the CIC. These operators implement the decisions of the SWC and TAO using the information entered by trackers and ID operators.

AAW decisions must be made very rapidly to be effective. As a result, the Navy has developed a philosophy of command by negation. This means that decisions are made by those directly involved in an operation — in this case they are made by AAW personnel at the unit level — and are reported to higher levels of command. The senior officers monitoring the decisions — typically the AAW coordinator and the force track coordinator — can either accept them or intervene and select a different course of action. This means that any of the unit AAW decisions listed above can also be made at the task force level, so aids designed for unit AAW decisions should also be capable of supporting higher levels of command. In addition, these aids should be designed to help officers on the task force staff monitor a large number of individual AAW activities and decide when to intervene.

The Characteristics of Unit AAW Decisions

The characteristics of AAW decisions at the unit level can be specified using the taxonomy in Table 1. A more complete version of this taxonomy, including a scale for assessing each characteristic is described in (7). In this section the taxonomies are used to assess the characteristics of unit AAW decisions. The assessments are those of the authors, based on interviews with Navy AAW personnel.

The first decision characteristic is the decision-maker's resources. For all unit level AAW decisions, the resources are those of a single ship augmented by any aircraft being controlled by the ship's CIC. In some situations two or more ships will operate together, especially when one has an automated AAW system and the other does not. In this case the TAO on one ship may control the AAW weapons on both, but usually the resources available to a ship are its own.

The decisions made during AAW operations are some of the most important encountered on a ship. The safety of ship, and, in some cases, the safety of the main units in the task force (e.g., an aircraft carrier or a troop carrier) depend on the decisions made by AAW personnel.

The number of alternatives available to unit AAW personnel depends on the type of decision. For example, there are basically two alternatives for threat detection: either a radar signal is identified as a track or it is not. On the other hand, a track can be identified as any of a large number of friendly or hostile aircraft and missiles. Similarly, many different assessments are possible for the danger posed by a threat and the capabilities of defensive weapons. Prioritizing the threats means selecting one of possibly many tracks to process next. The number of defensive weapons that are considered for assignment to a threat generally is limited to two or possibly three. There are a large number of options for the timing and conditions under which a weapon could be fired, but these are usually resolved by engaging as soon as possible.

All of the unit level decisions are influenced by a large number of factors, many of them uncertain at the time a decision must be made. A tremendous amount of information flows into a CIC during AAW operations on a variety of data and verbal channels. At every stage in the decision process, keeping track of all of the factors that potentially could affect or reverse a decision is a problem.

The number of attributes used to evaluate the outcomes of AAW decisions is low, in part because there is little time to balance conflicting objectives. This is especially true for the decisions associated with detection, tracking, and identification, where the accurate location and type of a track symbol is the predominant objective. The assessment and assignment activities can involve consideration of the need to defend the ship and maintain an adequate defensive posture in addition to defending the high value units in the task force. However, secondary objectives probably are ignored in combat situations involving many threats. Consideration of additional outcome attributes might be possible if automated aids were developed to support this activity, but this would require a more explicit statement of the relative importance of objectives than Navy personnel are currently willing to make.

In every decision activity, the relevant outcomes are defined in simple, measurable terms. For example, a track is either hostile, friendly, or unknown; and a defensive weapon is assumed to either destroy a threat or miss. Each decision-making function is defined with little ambiguity about its objectives and the extent to which they are achieved. Decisions requiring a subjective assessment of the relative value of various outcomes are referred to higher levels of command.

Contingent decision-making is primarily relevant for the assignment of defensive weapons to air threats and the specification of conditions under which they can be fired. Primary and backup weapons often are assigned at the same time with the implied decision that the second is to be used if the first fails. The other activity where contingent decisions can be relevant is in prioritizing threats, where a track may be given a low priority until it reaches a certain position or further information about it can be obtained. The other decision-making activities usually are made without consideration of future contingencies.

The outcomes relevant to all of the unit AAW decisions are dependent on a large number of uncertain factors. For instance, the accuracy with which a track is positioned or identified depends on the coverage and availability of radar systems, the effects of jamming, the weather, the proximity of friendly aircraft, the detection of electronic signals from the aircraft or missile, etc. Many of these factors are related. The need to make rapid decisions often makes it difficult for unit AAW personnel to consider the significance and implications of these relationships, with the result that many decisions are based on suppressing the dependencies and uncertainty inherent in the problem.

The degree of risk inherent in all unit AAW decisions is very high. The potential consequences of an improper decision include the loss of one or more ships, and the probability of this happening is sufficiently high to make the risk a dominant factor in the decision process. This risk is used as justification for suppressing all but the most significant uncertainties and objectives in reaching AAW decisions, although better decision aids could allow more thorough examination of each decision and potentially reduce the risk.

All unit AAW decisions are subject to review and approval, but, under the philosophy of command by negation, they are considered firm decisions until a higher level of command intervenes. If two or more units make conflicting estimates or actions, they may refer the decisions to higher authority for resolution, but most decisions are made under the assumption that they will be implemented.

The structure of current AAW decision activities has been well defined by the Navy, and all are conducted using standard procedures. While there are unique aspects to every AAW problem, the plans and procedures for dealing with them are specified in advance. This means that specialized decision aids can be designed for the AAW environment that do not require the user to formulate a new approach to the problem.

The quantity of information relevant to each of the decision activities processed in a ship's CIC during an engagement is extremely large. One of the major problems in designing any decision aid for unit AAW operations is to avoid overloading the CIC personnel with too much information. As a result, AAW decision aids must aggregate and simplify information wherever possible and provide detailed information only when needed. This does not mean that the aids should ignore significant information, especially subjective information about the credibility and accuracy of data sources.

Some of the pieces of information received by a ship's CIC are significantly more important than others. Most of the work required to filter incoming data and recognize important information currently is being done by CIC personnel. For instance, the TAO is expected to listen to several voice circuits simultaneously while watching a computer-generated display of tactical data and exchanging information verbally with other individuals in the CIC. Better decision aids could reduce this burden on CIC personnel and allow them to spend more time analyzing the decisions they face.

Most of the information sources used in the unit AAW operations are considered reliable, but CIC personnel often use crude rules of thumb to filter data from ships with inexperienced personnel or older equipment. The process of assessing the reliability or credibility of information sources is considered important by CIC personnel, but AAW decision aids do little to help these assessments. Thus, a great deal of emphasis is placed on communicating "accurate" information about tracks even when the data supporting that information is uncertain and incomplete.

The time available for all of the AAW decision activities is a few minutes or less during an engagement. Furthermore, the same types of decisions can be encountered repeatedly many times per minute in a major conflict. CIC personnel are trained to function in combat situations where the enemy attempts to saturate the air defenses of a task force. Obviously, decision aids designed for this environment must be capable of reacting with delays of at most a few seconds, a requirement that is not always met by existing systems.

5. DESIGNING IMPROVED AIDS FOR ANTIAIR WARFARE DECISIONS: AN APPLICATION OF THE TAXONOMIES

In the first part of this chapter we use the AAW decision characteristics described in the previous chapter as a basis for identifying the importance of each information processing and analytic function for AAW decisions. The importance of these functions is then compared to how well AAW personnel can perform them without an automated decision-aiding system. The functional needs identified in this manner are compared to capabilities of the Naval Tactical Data System (NTDS), and some of the functional limitations of NTDS are discussed. The last part of this chapter explores various functional improvements that could be made in existing AAW systems and describes ways that they might be implemented.

Functional Requirements for Unit AAW Decisions

The characteristics of AAW decision-making activities at the unit level provide a basis for identifying the more important information processing and analytic functions required to make each decision. These functions are listed in Table 3 and discussed in Chapter 3 of this report. Some of the decision functions were mentioned briefly as part of the discussion of decision characteristics in the previous chapter. This section compares the significance of these functions in making AAW decisions with the ability of unaided CIC personnel to perform them. This comparison leads to a set of functional requirements for AAW decision aids that can be used to evaluate existing decision aids and identify ways to improve them.

Although we do not have a precise way to measure the importance of an information processing or analytic function in a given decision activity, we are only concerned with identifying the relatively important ones. For this, a simple ranking of the functions based on a series of pairwise comparisons will suffice. The authors assessed such a ranking based on information in Navy warfare publications about AAW, and interviews over the past two years with a variety of Navy personnel who have experience with AAW operations, including several interviews with officers who train Navy crews to make AAW decisions at the unit level. We did not find unanimity among the Navy officers interviewed on the relative importance of individual decision functions, but their answers were sufficiently close to a consensus to identify the more important functions.

Our ranking of the importance of the information processing and analytic functions on a scale of 0 to 10 is shown in Table 4. The entries in the table are not precise, and differences of a couple of points on this scale are not significant. However, it is still possible to identify functions that are substantially more important than others for each decision in the table. Clusters of relatively important function/decision pairs are circled in Table 4. Each cluster is designated by a letter so it can be referenced in the following discussion.

The first significant information processing function is measuring physical data, such as radar signals (represented by cluster A in Table 4). The initial processing of this data takes place outside the CIC where the information is converted electronically to a form that can be relayed to the CIC and displayed on a radar repeater. The first human measurement occurs when the signal is displayed as a blip with a particular bearing and distance from the center of the screen. This kind of measurement is important for unit AAW activities concerned with detecting, identifying, and assessing potential threats, and with assessing the state of defensive aircraft and selecting the best time to fire defensive weapons. This function is well supported in a CIC, even on ships that do not have automated aids for AAW.

TABLE 4

THE IMPORTANCE OF INFORMATION PROCESSING AND ANALYTIC FUNCTIONS FOR AAW DECISION ACTIVITIES

(0=NOT IMPORTANT; 10=IMPORTANT)

		DETECT	I.D.	ASSESS	PRIORITY	COUNTER ASSESS	ASSIGN	SHOOT
ACQUIRE	MEASURE PHYSICAL DATA	8	8	8 A	4	8 A	3	8 A
INFORMATION	MONITOR COMMUNICATIONS	5	5	8	8	8 B	3	5
INTERPRET	ASSESS UNCERTAINTY	8	8	7	6	7 C	4	2
INFORMATION	ESTIMATE AGGREGATE DATA	2	2	10 D	3	10 D	2	2
STORE AND RECALL INFORMATION		7 E	5	8	8	8	8 E	2
	CATEGORIZE	2	2	6	10 F	6	5	2
STRUCTURE	SORT	0	0	1	10	1	5	2
INFORMATION	FILTER	6	6	6	2	4	5	2
	PATTERN RECOGNITION	9	8	7 G	3	4	3	2
	DEFINE ELEMENTS	2	2	2	2	2	3	2
ANALYZE	IDENTIFY TRENDS	5	5	8 H	3	3	3	2
DECISION	IDENTIFY RELATIONSHIPS	0	0	5	2	3	3	2
SITUATION	UPDATE RELATIONSHIPS	0	0	0	0	0	0	2
	NEW REPRESENTATIONS	2	2	3	1	1	1	2
	CALCULATE DATA	5	2	4	4	5	8 I	3
EVALUATE	EXTRAPOLATE OUTCOMES	1	1	3	2	3	10	3
DECISION	OPTIMIZE	1	1	1	1	1	10	3
SITUATION	SENSITIVITY ANALYSIS	1	1	1	1	1	8	3
	VALUE OF INFORMATION	1	1	1	1	1	6	3
	INTEGRATE RESULTS	4	4	4	5	4	4	1
TRANSMIT INFORMATION		8	8	8	8	8	10	8 J

IMPORTANT COMBINATIONS OF DECISIONS AND FUNCTIONS ARE CIRCLED

The second important function, monitoring communications, also deals with the acquisition of information. As shown by cluster B in Table 4, this function is especially important for assessing threats and the readiness of defensive weapons based on reports from aircraft and weapons operators under the CIC's control and communications with other commands. These communications are also important for prioritizing air threats. Extensive facilities exist in a CIC to support this type of communication.

Most of the information received by a CIC must be interpreted by AAW personnel before it can be used to reach a decision. Unfortunately, much of this interpretation is done intuitively in the final stages of the decision process when defensive weapons are assigned to targets rather than when the information is first received or generated. For instance, trackers and ID operators do not make explicit assessments of the uncertainty associated with a track, even though this kind of assessment can be important for subsequent decisions concerning the track (cluster C in Table 4). The trackers and ID operators often have to interpret incomplete and conflicting data, but they do so intuitively and do not routinely communicate their uncertainty to those who use the track information for subsequent decisions. This forces those prioritizing threats and assigning weapons to make their own assessments of uncertainty based on relatively little information. The uncertainty associated with assessing the danger posed by a threat and the capability of defensive weapons to deal with it is also important for subsequent decisions, but communicating this uncertainty is not as difficult because the same individuals are usually involved with the assessments and subsequent decisions. However, not using an explicit representation of uncertainty when assessing offensive and defensive weapons makes it difficult to communicate this information with other units.

Another form of interpreting information, estimating aggregate or summary descriptors from detailed data, is very important for assessing threats and defensive weapons (cluster D). In this case, CIC personnel are given a limited capability to express and communicate their estimates. Tracks can be classified as confirmed hostile, assumed hostile, unknown, etc., and this information is available to other decision-makers. However, decision aids do not help CIC personnel make or communicate explicit estimates of enemy capabilities and intentions during AAW operations. Similarly, it is important that AAW personnel estimate the readiness and capability of defensive weapons from data about the weapon's position, ammunition, damage status, etc., and communicate these estimates to those making allocation decisions. At present, this function is performed manually with little direct decision aiding other than the display of detailed data to the individuals making the estimate.

Storing and recalling information rapidly is important for most AAW decision-making activities (cluster E). This function is either done manually with vertical plots and status boards, or it is automated using computer data banks. In the latter case, manual procedures are used as a backup to guarantee that the information will be available even if a computer fails.

The principal information processing functions needed to prioritize threats are sorting and categorizing information about the threats (cluster F). For instance, it may be important to know which threats are closest to the high-value ships in the task force, and which are within the range of a defensive weapon. Some attempts have been made to help AAW personnel perform this function, but they have not given the user the flexibility to define criteria appropriate for sorting or categorizing threats in a specific combat situation. As a result, these functions tend to be performed intuitively without the direct support of decision aids.

Recognizing the existence of predefined patterns is the principal information processing function required to detect, identify, and assess threats (cluster G). Although some attempts have been made to provide decision aids for this function, most of the required information processing is done by CIC personnel. In situations where many tracks are being considered at once, the demands on unaided trackers and ID operators for pattern recognition may preclude them from spending an adequate amount of time on each track, or performing other functions such as assessing the implications of incomplete or uncertain data upon which the track is based.

A related function that is important for assessing threats is identifying trends in enemy tactics (cluster H). In this case, the pattern or trend is not predefined, but must be deduced from available data. Other than display the data, existing decision aids do little to help CIC personnel perform this function.

The most important analytic function required to assign defensive weapons to threats is the evaluation of alternatives. This function includes calculating the probable performance of alternative weapons, using a simple analysis to predict the outcomes associated with each alternative, determining the optimum alternative, and testing the sensitivity of the decision to changes in data or assumptions (cluster I). Existing decision aids perform some calcultions relevant to the assignment decision (e.g., intercept coordinates), but attempts to automate the functions mentioned above have not been very successful because the aids do not contain the judgemental interpretations of AAW information that are used by those assigning weapons. As a result, the calculations required to perform these functions are performed heuristically by CIC personnel.

The ability to transmit information from the CIC to weapons operators and other AAW command centers is very important for all AAW decision activities (cluster J). This function is well supported in a CIC.

Typical AAW personnel, working without automated decision aids, are better able to perform some of these information processing and analytic functions than others. For instance, humans can do a better job of recognizing patterns, especially if the data is displayed in an appropriate graphic format, than mentally storing and recalling large amounts of information. Given adequate communication channels, humans can do a good job of acquiring and transmitting information. In fact, one of the problems with existing automated AAW aids is that CIC personnel find it easier to communicate with other individuals verbally than with the aids, especially when the information being transmitted is qualified, subjective, or uncertain.

Table 5 contains an assessment on a scale of 1 to 10 of how well AAW personnel who do not have automated decision aids can perform the information processing and analytic functions required for AAW decisions. As in the previous table, the assessments are not intended to be precise, and they represent the opinions of the authors based on interviews with Navy AAW personnel. Although differences of a couple of points in these assessments are not significant, the range of estimates is sufficient to identify functions for which effective decision aids would be useful. The clusters of important function/decision pairs identified in Table 4 are reproduced in Table 5. Clusters in Table 5 that contain relatively low assessments (i.e., where unaided humans do not perform the corresponding decison function well) indicate a functional requirement for a decision aid. As shown in Table 5, the functions of acquiring and transmitting information are already performed well by CIC personnel, although there is room for improvement in the acquisition of information needed to assess the status and capabilities of threats and defensive weapons. All of the other important functions needed to make AAW decisions (i.e., the clusters C through I in Tables 4 and 5) can be viewed as a set of functional requirements for AAW decision aids.

In summary, these functional requirements are:

- assess and communicate the uncertainty and credibility of information produced by the detection, tracking, assessment, and priority setting activities;

- combine the information available about threats and defensive weapons into an aggregate estimate of capabilities and intentions;

TABLE 5
THE ABILITY OF HUMANS TO PERFORM INFORMATION PROCESSING
AND ANALYTIC FUNCTIONS FOR AAW DECISION ACTIVITIES

(0=LITTLE OR NO ABILITY; 10=AMPLE ABILITY)

		DETECT	I.D.	ASSESS	PRIORITY	COUNTER ASSESS	ASSIGN	SHOOT
ACQUIRE INFORMATION	MEASURE PHYSICAL DATA	8	8	6 A	8	6 A	8	8 A
	MONITOR COMMUNICATIONS	8	8	6	8	6 B	8	8
INTERPRET INFORMATION	ASSESS UNCERTAINTY	4	4	3	4	3 C	4	4
	ESTIMATE AGGREGATE DATA	2	2	2 D	2	2 D	2	2
STORE AND RECALL INFORMATION		3 E	3	2	2	2	2 E	3
STRUCTURE INFORMATION	CATEGORIZE	2	2	1	2 F	1	2	2
	SORT	2	2	1	2	1	2	2
	FILTER	2	2	1	2	1	2	2
	PATTERN RECOGNITION	5	5	4 G	5	4	5	5
ANALYZE DECISION SITUATION	DEFINE ELEMENTS	5	5	3	5	3	5	6
	IDENTIFY TRENDS IDENTIFY RELATIONSHIPS UPDATE RELATIONSHIPS NEW REPRESENTATIONS	8	8	6 H	8	6	8	8
EVALUATE DECISION SITUATION	CALCULATE DATA EXTRAPOLATE OUTCOMES OPTIMIZE SENSITIVITY ANALYSIS VALUE OF INFORMATION	4	4	1	2	1	3 I	4
	INTEGRATE RESULTS	5	5	2	3	2	4	5
TRANSMIT INFORMATION		8	8	8	8	8	8	9 J

IMPORTANT COMBINATIONS OF DECISIONS AND FUNCTIONS ARE CIRCLED

- store and recall subjective assessments of the status and capabilities of threats and defensive weapons;

- sort and categorize threats according to criteria appropriate to a specific combat situation;

- identify predefined patterns of data that indicate the existence and identity of tracks, and the capabilities of threats;

- identify trends in enemy tactics; and

- predict outcomes and evaluate alternative defensive actions.

It should be emphasized that performance of the last function in this list depends on adequate performance of the preceding functions.

The Functional Capabilities of Existing Decisions Aids: NTDS

Automated decision aids for unit AAW operations have been in use since the early 1960's. Most of these have been incorporated in the Naval Tactical Data System (NTDS), which continues to evolve as new equipment and procedures are developed and tested. Elements of NTDS provide some of the functional needs discussed in the previous section, with varying degrees of success. The experience the Navy has gained from using NTDS provides a starting point and guide for attempts to design and implement improved aids for unit AAW decision-making activities. This section explores the functional abilities and limitations of NTDS, especially those relevant to the functional needs discussed above.

NTDS is a specialized information processing system that accepts real-time data from radars, aircraft, weapons systems, operators, and other ships; processes this information and displays it on small-screen consoles; exchanges the information via digital data links with other NTDS units (or units with compatible data systems); and reports some of the information to other commands and non-NTDS units via teletype. The primary control and output device for this sytem is an NTDS console, which consists of a small screen resembling a radar repeater and a variety of control keys. There are several types of NTDS consoles, and most consoles can be operated in several different modes. However, the consoles use common data definitions and symbology, and information or instructions entered at one console can be observed at other consoles. For instance, a track detected using one ship's radar can be entered into an NTDS console and it will appear immediately on consoles aboard other ships with NTDS systems that are connected to the same digital data net.

The primary information displayed on an NTDS console is a set of "synthetic" symbols generated by the NTDS computers. Although radar signals can also be displayed directly on NTDS consoles, this raw data is not transmitted among NTDS units or used to allocate weapons. Information associated with NTDS symbols can be retrieved by console operators and displayed on panels near the screen. Digital communications through the NTDS computers is augmented by voice channels, both within a CIC and among ships and aircraft.

Information in NTDS is manipulated by numerous control keys and a moveable cursor on the screen. For instance, a track's speed and heading can be displayed by moving the cursor to the symbol on the screen representing the track and pushing the correct sequence of keys. The functions performed by the control keys change when the operating mode of a console is changed. The operating modes correspond to the difficult decision activities that take place in a CIC: detection and tracking, identification, weapons control, etc. NTDS consoles display a variety of symbols on the screen to represent friendly and hostile aircraft, missiles, ships, and submarines. The area covered by the display and the corresponding scale can be selected and adjusted by each console operator.

NTDS is expensive and installed on less than half of the Navy's ships. With the exception of some of the simpler tracking activities, NTDS operators need extensive training and practice to remain proficient. NTDS computers can process, store, and display a very large number of tracks, and they significantly increase the number of tracks that CIC personnel can monitor, control, and counter. Data transmissions between NTDS units are secure and capable of handing the necessary information flows. NTDS systems are complex and subject to failures, so they have been designed to be compatible with manual CIC procedures. In fact, a vertical plot and status boards usually are maintained in parallel with NTDS operations in case the NTDS fails.

NTDS does not acquire the information needed for AAW activities or determine its accuracy. As raw information is received from a ship's radars, it is displayed to CIC personnel, but NTDS does not process, communicate, or store it. The only information that NTDS processes is that associated with the synthetic video symbols. This information is usually entered by detector/trackers and ID operators, although some automatic or semiautomatic tracking devices have been used in conjunction with NTDS. Semiautomatic tracking requires an operator to detect and initiate each track; manual tracking with NTDS also requires the user to reposition the synthetic symbols whenever necessary to follow a radar signal.

Translating radar signals into information that can be processed by NTDS requires a sequence of subjective judgements that are neither supported nor adequately documented by NTDS. The process is very tedious and is performed by enlisted men who are trained on the job. <u>However, the quality of all subsequent decisions depends on their judgements.</u> NTDS provides no way for them to communicate the confidence they place in their judgements, and no way for others to review at a later time the raw data from which the judgements were made. Inaccurate data and incorrect interpretations of incomplete data are rapidly disseminated to other NTDS users and the source of an error often is difficult to identify.

The solution to this problem is not to store and recall large amounts of raw data in NTDS. The reduction of this data to a relatively simple description of the tactical situation is essential for subsequent AAW decisions. However, automated aids to support this kind of data reduction should provide a means for communicating and processing the credibility and uncertainty of the resulting information. Artificially eliminating uncertainties by establishing a symbol on an NTDS screen or refering conflicting opinions to higher levels of command does not provide tactical AAW decision-makers with an adequate picture of reality. The challenge is to find a way to represent uncertainty and credibility in a way that is easy to communicate and interpret.

NTDS does not accept or process intelligence estimates or other prior judgements about enemy tactics or intentions even though this information can be useful for detecting, identifying, and assessing air threats. All processing and updating of subjective estimates must be done by AAW personnel. Some of this processing could be done by an automated decision aid if it were capable of accepting and communicating estimates of uncertainty and credibility. This capability would also allow an aid to help a user recognize anticipated patterns in enemy tactics.

The process of entering information and instructions into an NTDS console is slow relative to verbal communications. Communicating rapidly via a set of function keys is not a natural activity for most individuals, especially in a crisis. As a result. information often reaches CIC personnel via voice circuits before it shows up on an NTDS console. In rapidly developing situations, NTDS can resemble a ticker tape that is running late. In this case, there is an even greater incentive to rely primarily on the voice circuits that are supposed to augment NTDS rather than compete with it. Discipline is required to see that operators use NTDS as their primary communications channel. A simplified method of entering data and instructions into NTDS oriented more toward the user than the computer could help to overcome this problem.

However, even when all the required information is stored in NTDS, it is not always used. The complexity of NTDS controls and displays cause many users to overlook information relevant to their AAW activities, often because they do not know how to access it. The same is true of some information processing features available in NTDS. Additional training of CIC personnel would help to overcome this problem, but a better solution would be to simplify the NTDS displays and controls. NTDS operators currently have limited control over the amount and type of information displayed on the screen. For instance, they can choose to eliminate subsets of the tracks from the display (e.g., all tracks representing ships and submarines). However, the displays are still very complicated, and NTDS does not help the user aggregate or interpret the information displayed. In addition, NTDS users are faced with the task of integrating a great deal of information arriving simultaneously on the screen, several other NTDS display panels, and a variety of voice circuits.

Attempts to incorporate decision algorithms in NTDS to evaluate alternative weapons allocations have not been successful. Most NTDS computers are programmed to allocate weapons automatically or make recommended allocations, if the user chooses to have the system perform these tasks. However, most AAW personnel consider this system for evaluating and optimizing alternatives to be worse than useless. They believe the system does a poor job of allocating weapons and distracts them when they are trying to make their own decisions. The problem is not that the allocation system is faulty, but that it is an oversimplified algorithm based on only part of the information available to AAW personnel. A prerequisite for designing an adequate decision algorithm is to give NTDS the capability of processing all of the significant information used by human decision-makers, including subjective judgements about the validity and uncertainty of track data. An improved algorithm also should allow the user to interact with the system to test the effects of different decision criteria and assumptions.

In summary, NTDS is primarily a passive information processing and communication system. It automates many of the routine information processing activities conducted in manual AAW operations (e.g., plotting tracks, computing intercepts, communicating estimates and decisions) with the result that CIC personnel can perform these tasks more rapidly and for a larger number of tracks. NTDS is very good at accepting, storing, and displaying data in certain standard formats, but there is little flexibility to process information that cannot be expressed using these formats. Also entering and assessing data in NTDS is sufficiently slow and complex that important information can be ignored or overlooked in rapidly developing combat situations. NTDS does little to integrate, analyze, or evaluate AAW information. Instead, it displays basic information to the AAW personnel who perform these functions.

Improved Aids for AAW Decision-Making Activities

The analytic procedures and taxonomies described in this report can be used to establish a set of functional needs or goals that guide the development of improved decision aids. However, they do not determine exactly how the aids should be designed. In fact, there may be several designs that provide some or all of the required functions. Research and development is needed to design and test aids that can provide the functions required for unit AAW decisions. This section indicates some of the directions in which this design effort might proceed. The designs discussed here are conceptual and preliminary; they are intended as a starting point for developing better AAW aids.

Probably the most fundamental improvement that could be made in AAW decision aids is to provide them with the capability of accepting, processing, and displaying uncertainty. Many of the other functional improvements depend on implementation of this capability. Developing aids that can do the calculations needed to provide this capability is not likely to be a major problem even on limited computer systems. There are ample numerical and analytical techniques for reducing the calculations to a manageable level. The difficulty lies in defining simple ways to enter and display uncertainties that will not require too much time of AAW personnel or overwhelm them with too much information. Conventional methods of expressing and displaying uncertainty (i.e. probability distributions) are not appropriate given the volume of data and time available for AAW operations.

There are several ways to solve this problem. For instance, point estimates of scalar quantities (e.g. the number of enemy aircraft in an attacking group, or the time required for a defensive weapons system to rearm or get in position) can be accompanied by a range of likely values. Similarly, uncertainty concerning a track's position can be estimated as a region where the track is likely to be, similar to the "circular error probability" used to assess uncertainty in strategic weapons, or the elliptical probability contours used in ASW command and control. A representative display of this kind of information is shown in Figure 2. Uncertainty about the existence of a threat (e.g. the possibility that a low flying threat has not been detected) could be represented by shaded regions or probability contours indicating areas where radar coverage is less likely to detect certain types of threats. Figure 3 shows one possible way to display this information. Uncertainty about the identity of a threat could be indicated by the shape or color of the track symbol. (This is one area where NTDS currently provides some means of expressing uncertainty.) None of these approaches provide a complete description of the uncertainty involved, but they would allow a much more accurate description of a tactical AAW situation than the artificially deterministic estimates in current AAW systems.

Would the inclusion of uncertainty in AAW data overwhelm decision-makers with too much information? There is already a potentially overwhelming amount of data in NTDS, and yet AAW personnel still find it necessary to consider and mentally process uncertainties in order to reach a decision. If some of this processing were done by a decision aid and displayed in a simple format, the decision-maker could devote his attention to

FIGURE 2

DISPLAYING THE UNCERTAINTY IN A TRACK'S POSITION

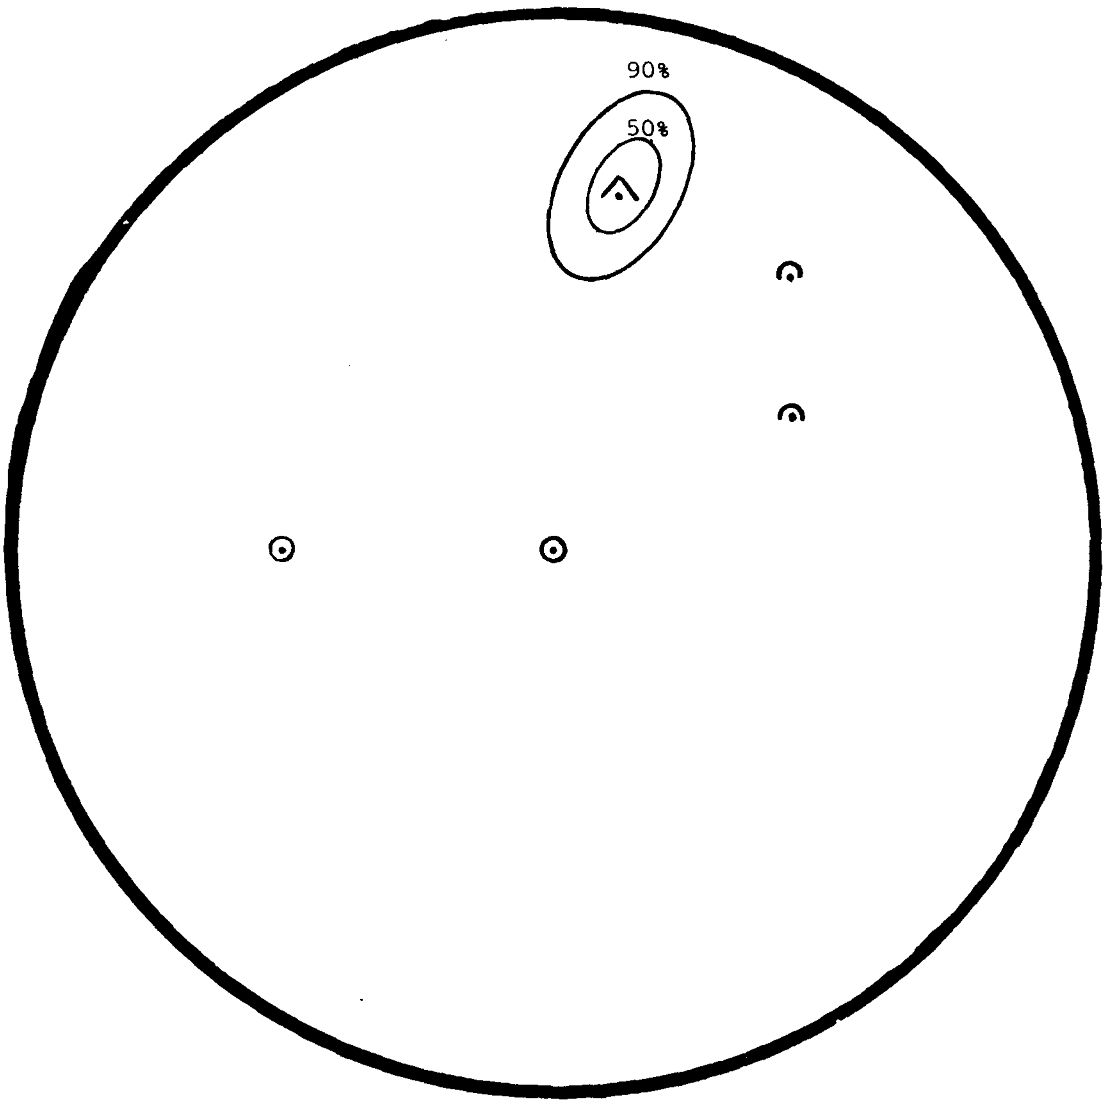

There is a 50% chance that the aircraft is in the inner oval, and a 90% chance that it is in the outer oval.

FIGURE 3

DISPLAYING UNCERTAINTY ABOUT THE EXISTENCE OF A THREAT

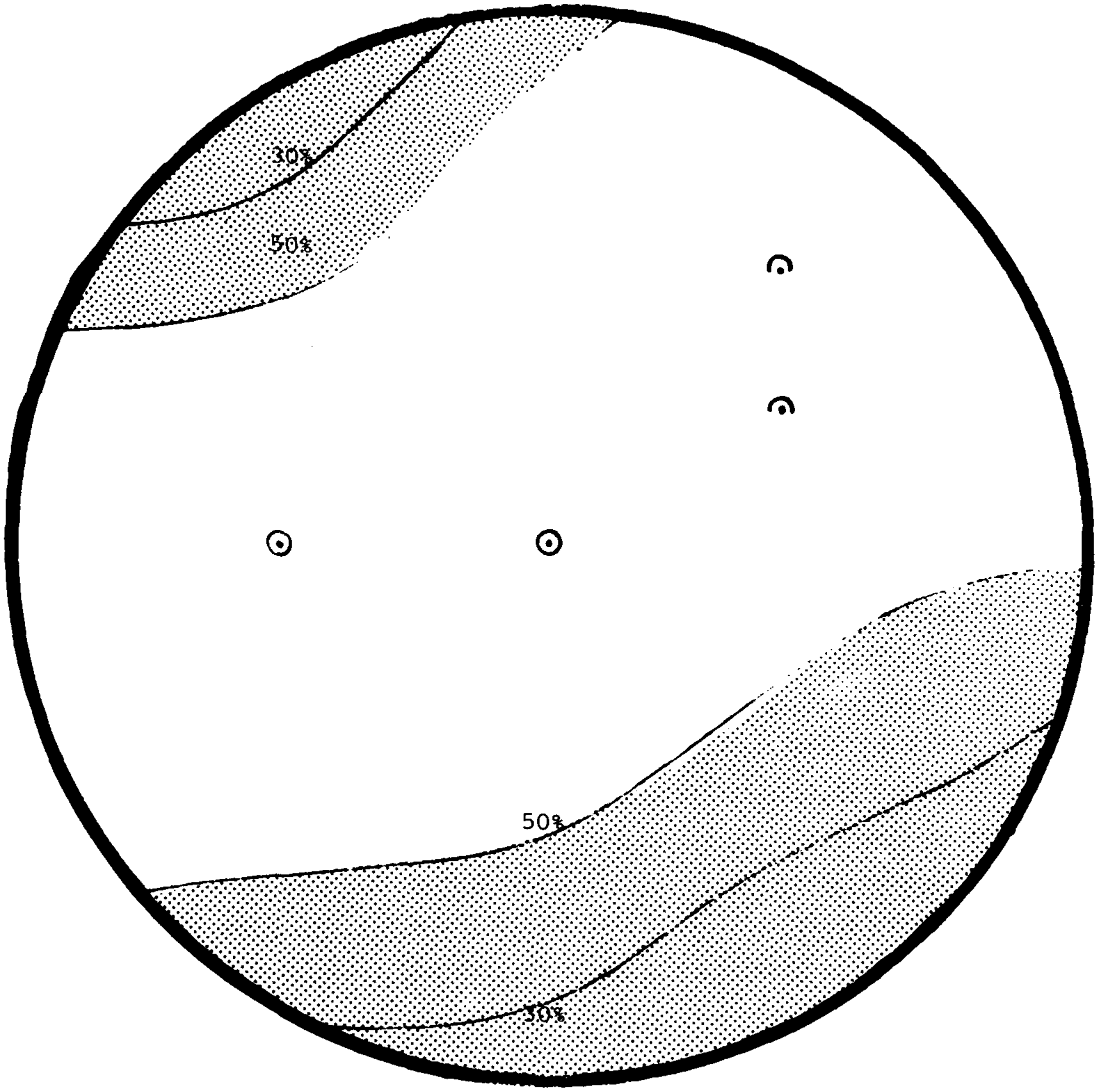

The shaded regions are areas where the probability of detecting a low-flying threat is less than 50%. Beyond the second set of contours, the probability of detection is less than 30%.

interpreting the results rather than trying to calculate them. For instance, an automated display of areas where undetected threats are more likely to exist would indicate the value of holding defensive weapons in reserve as a hedge against the sudden appearance of a threat near the task force. Existing decision aids do not burden AAW personnel with this information, but they also provide little guidance for dealing with undetected threats.

By processing uncertain data and estimates, an aid can measure and display the relative credibility of various pieces of information. The logic required to process the credibility of data is not so well developed as probability theory, but the two are closely related. The basic elements of credibility analysis are outlined in Chapter 8 of this report. An aid incorporating credibility algorithms could supply the user with an indication of both the uncertainty in AAW data and the extent to which the data is based on reliable sources of information and analysis.

An aid capable of processing uncertain data could play a more significant role in detecting, tracking, and identifying aircraft and missiles. By matching data received about tracks with the known characteristics of various aircraft and missiles, an automated aid could calculate the likelihood that the data is caused by various offensive and defensive weapons systems, even if some of the data is missing or contradictory. The same type of information processing would indicate the probability that data was generated by <u>any</u> aircraft or missile, thus helping the user detect tracks from noisy data. A track's position could also be updated probabilistically as new information about it is received, even if radar signals are not received continuously. A probabilistic calculation could automatically and continuously transform the track corresponding to a lost radar contact into a region where an undetected track is likely. See Figure 4.

FIGURE 4

TRANSFORMING A LOST TRACK INTO A REGION WHERE AN UNDETECTED AIRCRAFT IS LIKELY

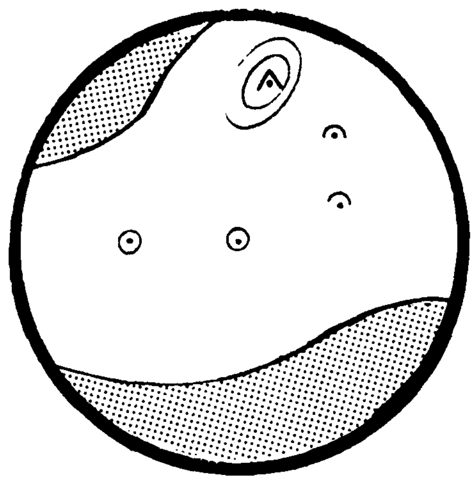

A) DISPLAY PRIOR TO LOSING RADAR SIGNAL

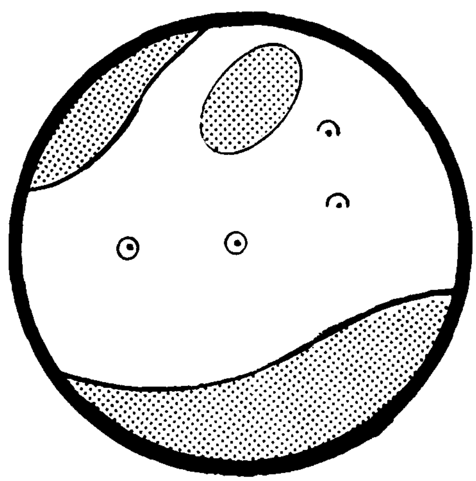

B) DISPLAY AFTER LOSING RADAR SIGNAL

Shaded regions are areas where an undetected threat is likely.

The algorithms required to do this sort of processing must be capable of testing several hypotheses for the origin of observed data, including "none of the above". In other words, the aid should determine the likelihood that a threat is not present, as well as the likelihood that a particular threat is near the task force. Methods exist for drawing these sorts of inferences by updating nonparametric relationships between observed data and the possible causes.

Once AAW decision aids have the capability of processing subjective estimates of uncertainty, they can be given the flexibility to accept and use prior estimates of enemy tactics and intentions from intelligence sources or CIC personnel. For instance, a Bayesian pattern recognition algorithm would allow users to input estimates of the most likely timing, location, and type of enemy attacks (e.g., a low altitude attack from the north). See Figure 5. The system would then be more sensitive than usual to radar data indicating the anticipated attack. The same kind of processing would use estimates of the credibility of various information sources to govern their influence on the data being updated in the system. Probabilistic calculations can also be used to identify areas where additional intelligence or better surveillance would be especially useful.

It is difficult to make probabilistic intelligence estimates of events that can occur at any time and in any order because the timing and likelihood of each event may depend on the occurence of others. Direct estimation of the probabilities of all combinations and sequences of events is prohibitively difficult and time consuming. However simplified techniques for estimating approximate event probabilities and timing have been developed. These techniques could be adapted to the AAW environment to provide a means of helping AAW personnel incorporate intelligence information into their decision-making activies.

FIGURE 5

DISPLAYING INTELLIGENCE ESTIMATES OF THE DIRECTION AND TYPE OF ENEMY AIR ATTACK

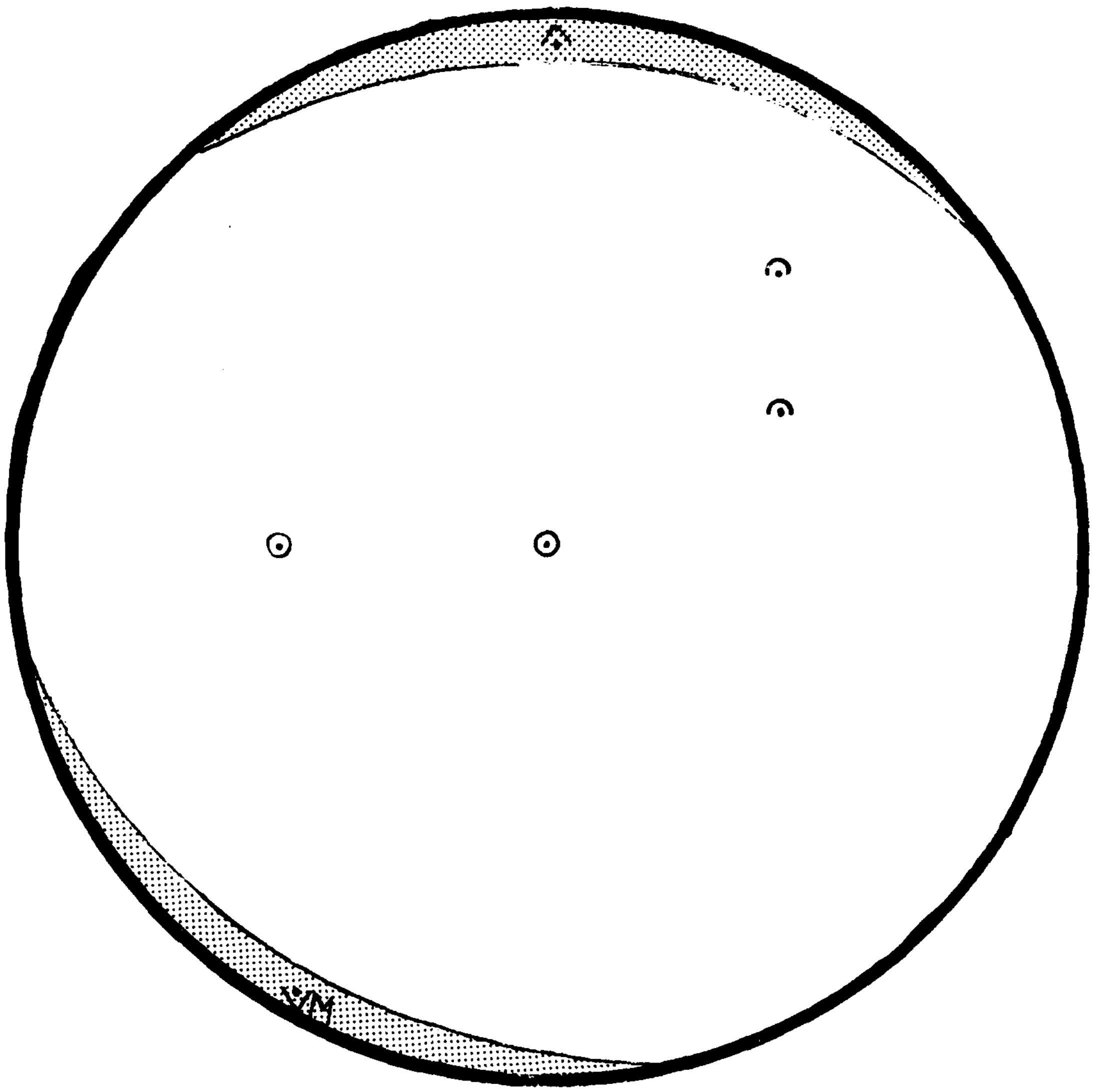

The shaded borders indicate directions from which an attack is likely. The thickness of the border is proportional to the probability of an attack from that direction. The symbols in the border indicate that an attack by aircraft is expected from the north and a submarine-launched missile attack is expected from the southwest.

Another major area where AAW decision aids could be improved is the generation of better summaries of the tactical situation. NTDS does very little to aggregate or summarize the detailed information it contains, forcing users to do the necessary information processing. An automated aid that summarizes AAW information should display the result graphically to help user interpret it quickly. Several different displays would be appropriate for summarizing different aspects of AAW operations.

One display consists of probability contours indicating the likelihood that threats will be detected by available radar and surveillance units at any location near the task force, or the likelihood that defensive weapons can intercept and destroy them at any given position. See Figures 3 and 6. These displays would summarize visually the status and capabilities of defensive surveillance and weapons systems, based on their performance parameters and reports about casualties, equipment failures, etc. The displays would make it easy to identify quickly the existence and size of gaps in the air defenses of a ship or task force.

A related display, which would provide a quick assessment of air defense capabilities, consists of a set of probability contours showing the likelihood that hostile aircraft or missiles could penetrate to any position in the vicinity of the task force. See Figure 7. These contours could be calculated from current position and status of threats and defensive weapons together with a simple model of the likelihood that the weapons will be able to counter each threat.

Another display could summarize the development of an air engagement over time to help AAW personnel identify trends in enemy tactics. Existing AAW decision aids are designed to give users a picture of the current situation, but do little to help them relate this information to earlier developments. The task of remembering earlier enemy tactics and performance, and drawing inferences for current AAW decisions, is done by humans, even though the

FIGURE 6

DISPLAYING THE LIKELIHOOD THAT DEFENSIVE WEAPONS CAN DESTROY A THREAT

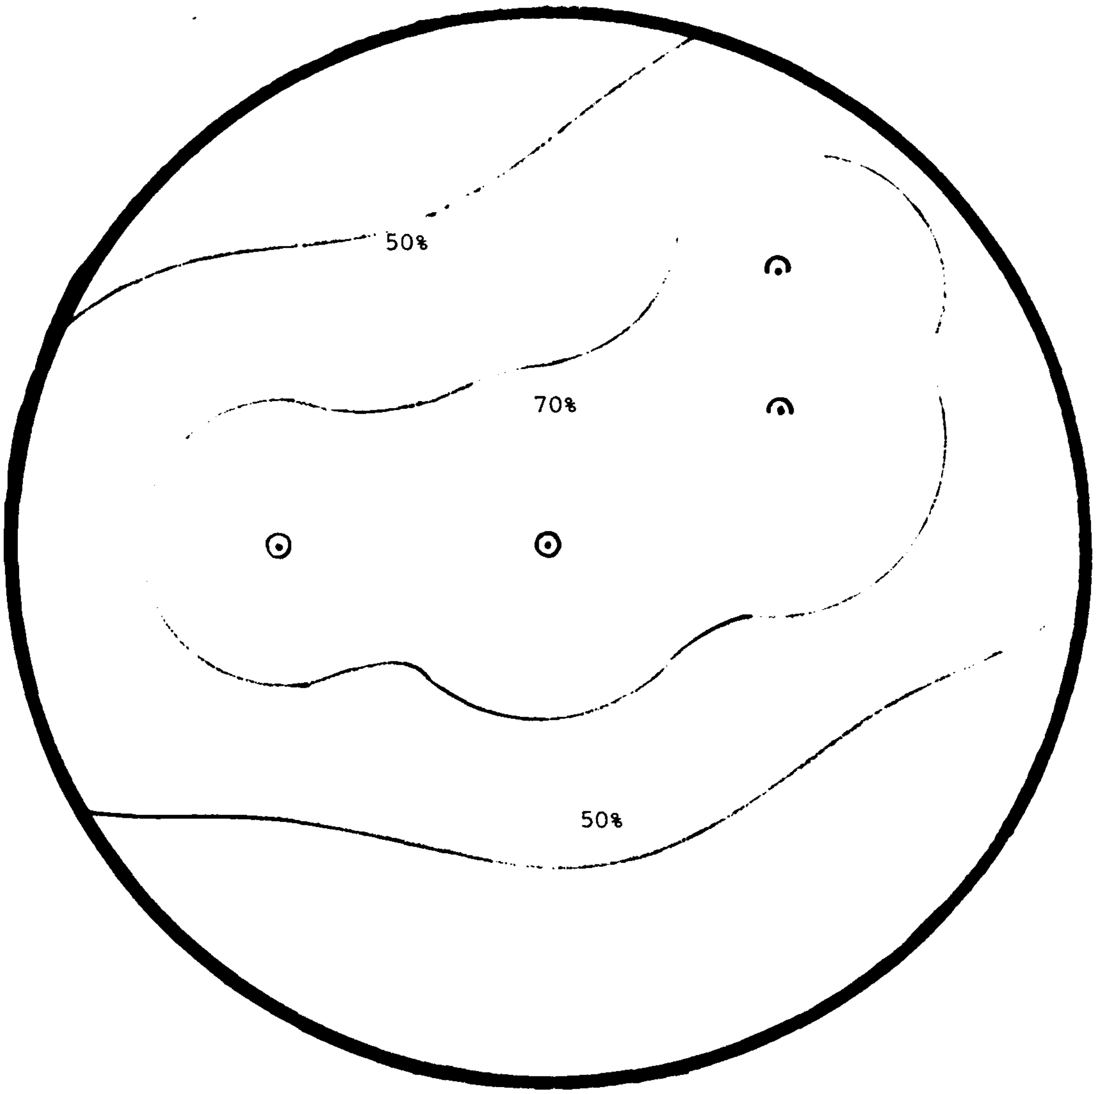

The probability that an air threat can be destroyed if it is inside the inner contour is greater than 70%. The probability is less than 50% outside the outer contour.

FIGURE 7

DISPLAYING THE LIKELIHOOD THAT AN AIR THREAT CAN PENETRATE TO ANY LOCATION

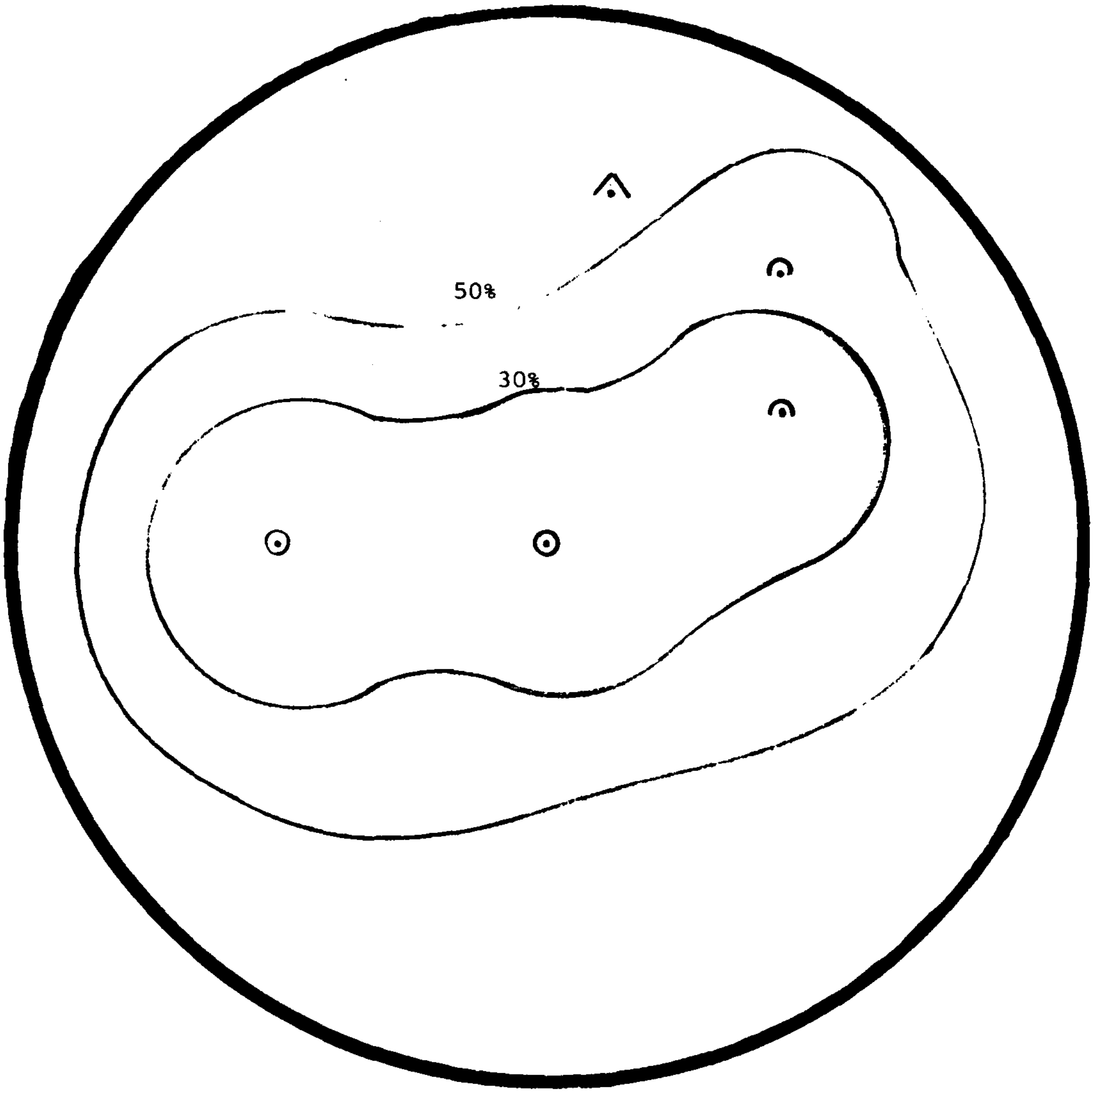

The probability that an air threat can penetrate to position inside the inner contour is less than 30%. The probability it can reach a point outside the outer contour is greater than 50%. This display is based on the likelihoods of both detection and destruction of air threats.

information needed to support this task has already been entered into an automated aid. A display showing the past movements of threats (perhaps with time markers along each track) or shaded areas indicating regions of heavy enemy activity would simplify this task and make it harder to overlook important aspects of enemy activities. One possible summary of enemy activity is shown in Figure 8.

A more significant summary of enemy tactics could be generated by identifying enemy movements or actions that are related. For instance, it may not be obvious from a large number of aircraft movements during the course of a battle that low-level attacks are usually preceded by a high-altitude feint from another direction. However an automatic aid could aggregate individual pieces of data associated with enemy actions and identify broader trends. This would help AAW personnel improve defensive tactics and resource allocations, as well as summarize the status of an engagement.

One area where improvements have been attempted in existing aids is the evaluation of alternative allocations of defensive weapons. These attempts use simple algorithms to recommend or make weapon assignments, but do not allow the user much interaction with the evaluation process. Once aids have been designed to accept subjective estimates, it will be possible to design much more flexible allocation algorithms. For instance, the aid could allow the user to specify various criteria for prioritizing threats and then sort or categorize the threats accordingly. Since the user's criteria may change rapidly during the course of an air battle, the aid must provide a simple way for defining criteria, including a menu of standard criteria (e.g., minimize threat to the task force, minimize threat to own ship, maximize expected number of missiles destroyed, etc.). By combining or modifying these criteria, the user could work interactively with the aid to prioritize threats and allocate weapons.

FIGURE 8

DISPLAYING A SUMMARY OF AN AIR ENGAGEMENT OVER TIME

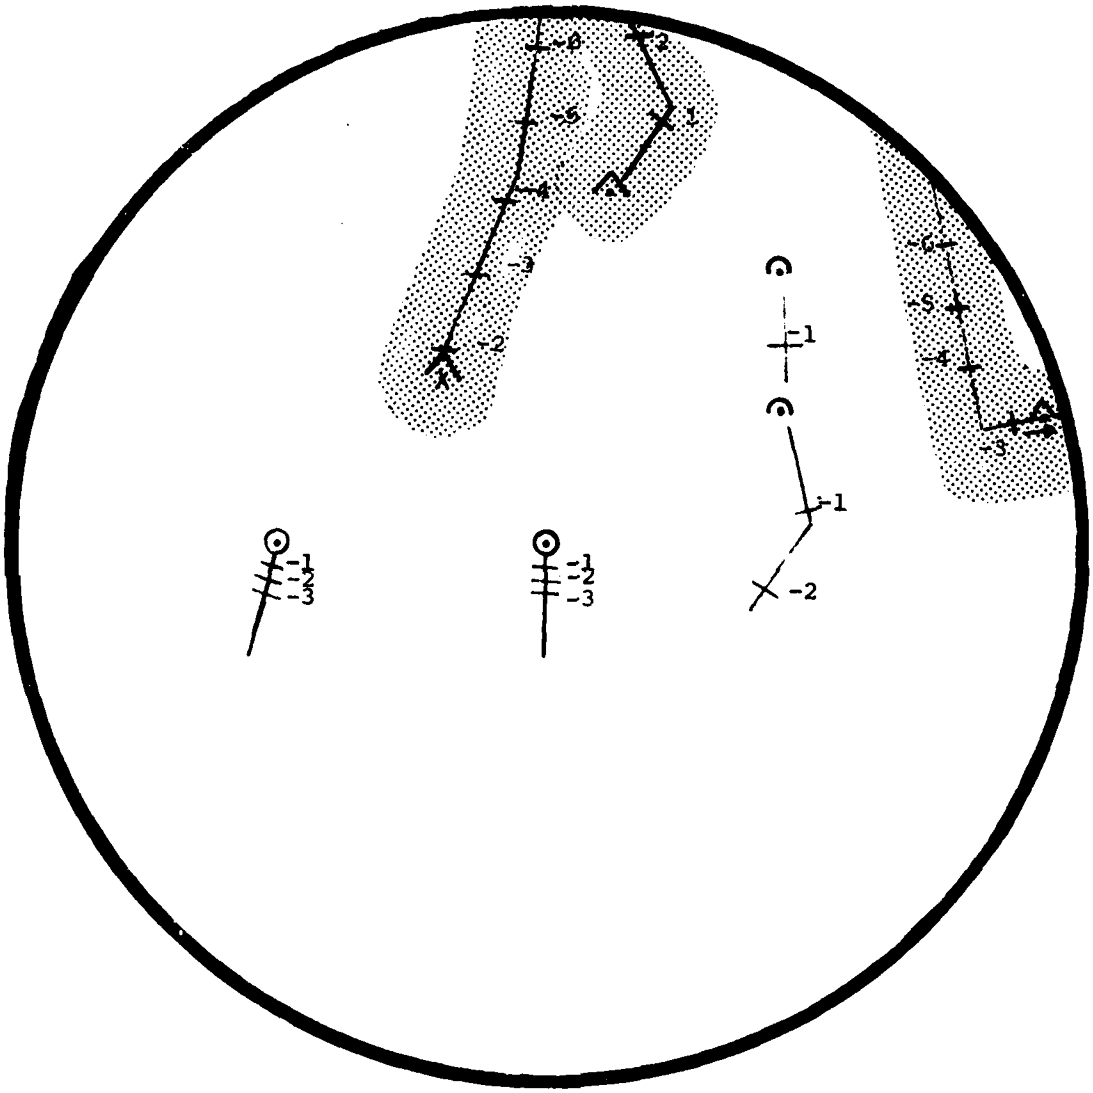

The lines indicate paths of friendly and hostile units. The markers on the paths show positions at earlier times. The shaded areas are close to positions recently occupied by enemy aircraft. The symbol with an "X" indicates an enemy aircraft that was destroyed. The symbol at the right edge of the screen indicates an enemy aircraft that retreated off the screen.

A similar interactive facility could be provided for predicting the outcomes of air engagements using simple models of the effectiveness of alternative defensive weapons against various threats. These evaluations, based on the kind of subjective information discussed above, would identify the risks associated with alternative defensive actions and suggest additional moves to compensate for them. After tentatively selecting an allocation of defensive weapons, the user could use the system to test the sensitivity of the decision to changes in assumptions (e.g., would a defensive allocation leave the ship in a vulnerable position if another attack occurs). The effects of an allocation of defensive weapons could be summarized in a display like the one shown in Figure 7.

There are many possibilities for improved aids to evaluate AAW resource allocations. However, their design depends on other improvements in the information processing and analytic abilities of AAW decision aids. Attempts to design better resource allocation aids are likely to meet the same fate as current systems if they are not preceded by improvements in an aid's ability to interact with the decision-making process, and capture in a simple but realistic way the uncertainties and subjective information relevant to his decision.

PART II

THEORETICAL EXTENSIONS OF THE TAXONOMIES

6. THE ROLES OF TAXONOMIES AND ANALYTICAL PROCEDURES IN DEVELOPING AND EVALUATING DECISION AIDS

One of the objectives of this research is to determine whether taxonomies can be used meaningfully to evaluate and guide the development of existing and proposed decision aids for Navy tactical command and control. In the previous phase of this research, we developed several related taxonomies to help someone think logically about the many aspects of decision situations and decision aids. This report describes an additional taxonomy of the functions required of an aid or decision-maker to analyze and make decisions, and a process for using the taxonomies to arrive at a logical framework for describing appropriate aids for various classes of decisions.

However, in the course of our research we found the need for more precise analytical procedures than taxonomies to evaluate at least some of the characteristics of decision aids. While these procedures are more difficult to develop than the taxonomies, they are potentially better measures of the value of decision aids. Several of these analytical approaches are discussed in this report. They are not fully developed, and deal with only a subset of the issues addressed in the taxonomies. However, they have been used to clarify concepts used in the taxonomies, and they show promise for augmenting or replacing portions of the taxonomies.

Taxonomies represent a general conceptual approach for thinking about decision aids. They have several important advantages that guarantee that they will remain useful even after more precise evaluation procedures have been developed. Taxonomies provide a comprehensive checklist to ensure that all relevant aspects of an aid and the decisions it supports are considered. They can also organize the evaluation process into a series of logical steps by showing the sequence in which issues should be considered and compared. Taxonomies also provide a framework for decomposing general characteristics of decision situations and aids into basic elements that can be assessed or measured. This decomposition provides a guide for establishing the relative importance of desired characteristics. Finally, one of the major strengths of taxonomies is that they do not require precise, unambiguous assessments of the characteristics they contain, thus making it easy to use them.

However, the ambiguity permitted by a taxonomy is also one of its major weaknesses. It may not be possible to interpret correctly imprecise assessments of an aid's characteristics or capabilities, or to determine the implications of these assessments for the applicability of the aid in various decision situations. For instance, what does it mean for the usefulness of an aid if someone says an aid is "reliable" or a decision is "risky"? Taxonomies can provide a guide to evaluating aids, but not an operational measure of their value. We need an explicit, reproducible, and logically sound method for processing the information derived from a taxonomy to determine how much should be spent developing a particular aid.

The ambiguity inherent in taxonomies also means that they are not unique. Different taxonomies can be used to decompose and organize the same characteristics of a decision or an aid. The adequacy of a particular taxonomy is a matter of judgement.

While taxonomies provide a useful and flexible framework for conducting a general evaluation of decision aids, they should be augmented by more precise analytical procedures to guide specific resource allocations. Unfortunately, the analytical procedures and concepts needed to evaluate decision aids are not well developed. For instance, we do not have a basic theory to describe the value of decision aiding, analysis, and modeling, although several attempts have been made to develop one. (A survey of some of the work in this area is outlined in Chapter 7 of this report.) In addition, there is no formal measure of the credibility of decision aids, although it is clear that an aid must be credible to be useful to a decision-maker. (The issue of credibility is discussed in detail in Chapter 8.) Analytical procedures such as these are at least as difficult to develop as taxonomies, but they provide an unambiguous and reproduceable way to evaluate at least some aspects of decision aids. Both taxonomies and more formal analytical procedures rely on subjective assessments of the properties of decision situations and aids, but taxonomies also require a subjective integration of these assessments to evaluate or design an aid. Analytical procedures rely on an explicit and testable set of logic rather than the intuition of the evaluator.

We anticipate that taxonomies and analytical procedures eventually will complement each other, and will be used together to evaluate decision aids. An evaluation procedure using both would probably start with taxonomies to make sure the important issues or features associated with an aid are considered and to identify the key tradeoffs that must be made. Then one or more rigorous analytical procedures would be used to produce formal measures of an aid's potential value. Finally, general evaluation procedures based on taxonomies would be employed to deal with aspects of an aid or decision situation that are poorly defined or for which no formal evaluation procedure exists. The question is not whether it is best to use taxonomies or more formal analytical procedures, but rather how we can integrate the results of both approaches. (See Chapter 8 of this report for the initial elements of a theory dealing with the use of formal analysis to update more intuitive logic.)

7. A REVIEW OF RESEARCH ON LIMITS TO HUMAN INFORMATION PROCESSING AND DECISION ANALYTIC ABILITIES

The procedure for developing and evaluating decision aids using a taxonomy of decision functions is based on the idea that humans have limited abilities to perform certain functions, and aids may be able to overcome these functional deficiencies. In order to understand how an aid can help the decision-making process, we must first understand how humans make decisions without an aid. This chapter reviews some of the research that has been done on the limits to human rationality, and the extent to which decision aids and analytic procedures can overcome these limits.

The organization of this chapter is as follows. We first review several alternative theories of how humans make decisions. These theories are based on an interpretation of alternate forms of decision-making behavior as manifestations of different bounds on rationality. These bounds may be viewed as scarce resource constraints that influence the extent to which a decision-maker can act rationally in a given situation. We then discuss the degree to which analytical methods and decision aids can or should reduce these bounds to rationality, and the extent to which these methods introduce undesirable attributes of their own.

There is a great deal of evidence that unaided humans often make very poor decision-makers by any measure. Empirical research by psychologists, organizational behaviorists, decision analysts, and others has revealed that in many situations humans tend to be poor information processors, poor at computation, biased evaluators, and inconsistent at making choices, although some individuals seem to be naturally good decision-makers. Over the last twenty five years a great deal of progress has been made in understanding why people behave the way they do. Although a complete review of the literature is beyond the scope of this report, the work of three authors — Simon, March, and Radner — provides a good framework for this discussion.

Simon starts from the idea that all intendedly rational behavior is behavior within constraints (15,16). In his view the most important of these constraints reflect limitations in human computational capabilities and the organization and utilization of human memory. Simon makes two relevant points for this discussion: first, that the state of information of a decision-maker is an important characteristic of his decision process; and second, that a decision-maker may deliberately introduce simplification into a model of a situation in order to bring the model within the range of computational capacity. Simon concludes that a theory of rational behavior must be as concerned with rational actors coping with uncertainty and cognitive complexity as with the objective environment in which they make their decisions.

Simon distinguishes between two different types of rationality: <u>substantive rationality</u>, the extent to which appropriate courses of action are taken; and <u>procedural rationality</u>, the effectiveness, of the procedures used to choose these actions in light of a decision-maker's cognitive powers and limitations. Conceptually, there is no reason to treat the substantive and procedural problems as though they are independent of each other: The global optimization problem facing a decision-maker is to find the least cost or best return decision, including computational cost. In practice, however, this is rarely the way humans proceed. Not only are there cost constraints on the computation of solutions of problems, there are also constraints on time, memory, imagination and a variety of other phenomena. These constraints are difficult to overcome and even more difficult to take into account in a systematic way.

Radner also deals with the problems of what he calls "limited rationality". He characterizes a number of different types of behavior that seem to follow from a combination of different types of bounded rationality: (1) constant proportions, in which the allocation of effort is constant over time; (2) putting out fires, in which all effort is allocated to those activities that have the worst performance at any given time, and (3) staying with a winner, in which all effort is allocated to those activities that have the best performance at a given time. Radner emphasizes that such behavior is not necessarily the result of costly information and analysis, but is also influenced by the "limited capacities of humans and machines for imagination and computation." Since formal methods will never take the place of imagination, it is clear that we will always have to live with some form of bounded rationality.

March provides another more detailed characterization of types of rationality. He distinguishes between calculated and systemic rationalities (5). Action, he says, is presumed to follow either from explicit calculation of its consequences in terms of objectives, or from rules of behavior that have evolved through processes that are sensible but that obscure from present knowledge full information on the rational justification of any specific rule.

The distinction between calculated and systemic rationalities is a good starting point for discussing the role of decision aids. Four of the calculated rationalities identified by March are described below (the definitions are taken from (5)). Identifying the implicit resource constraints associated with these rationalities provides a natural framework for evaluating the potential impacts of analysis.

Limited rationality emphasizes the extent to which individuals or groups simplify a decision problem because of the difficulties of considering all the alternatives and information. This type of rationality manifests itself in oversimplified decision procedures such as heuristic search rules, incremental thinking, muddling through, and uncertainty avoidance. A scarce resource here is the ability to deal with complexity.

Contextual rationality emphasizes the extent to which choice behavior is embedded in a complex of other claims on the attention of decision-makers. The scarce resource that is emphasized is time: When the decision-maker cannot spend much time on a given decision, it is likely to look "irrational" when compared to other more carefully considered decisions.

Process rationality emphasizes the extent to which decisions find their sense in attributes of the decision process rather than in attributes of the decision outcomes. Such behavior is typical in large organizations in which, for purposes of control, many decision procedures are standardized. Decision-makers are rewarded not for good outcomes, but for adherence to the "rules". Flexibility and control may be viewed as scarce resources here.

Game rationality emphasizes the extent to which organizations and other social institutions consist of individuals who pursue individual objectives by means of calculations of self interest. The scarce resource here is cooperation.

For each of these types of rationalities a case can be made that decision aids can be of help in removing the bounds, or at least in providing a way of dealing with them. For example, an aid based on decision analysis could help someone deal explicitly with complexity and uncertainty, and enable them to do quick computations, even in very intricate situations. This type of aid is focused on reducing the bounds inherent in limited and contextual rationality. Similarly, an aid to modeling organizational behavior would focus on the bounds of process rationality by providing a framework for taking into account procedural or organizational problems explicitly in the decision process. An aid employing game theory addresses game rationality by offering a structure for identifying and analyzing ways in which cooperation can help the individuals in a group all achieve better outcomes. Thus, decision aids have the potential to influence positively the factors inherent in the calculated rationalities. The scarce resources are relatively easy to identify, and also relatively easy to influence.

On the other hand, the case for decision aids is somewhat harder to make for the systemic rationalities. Here the scarce resources at play are somewhat harder to identify and also more difficult to influence. March identifies three types:

Adaptive rationality emphasizes experiential learning by individuals or organizations that permits the efficient management of considerable experiential information, but tends to separate current reasons from current actions. This type of rationality describes the intuitive decision-maker who has a hard time justifying his choices to others. The scarce resource here seems to be the ability to understand and articulate the reasoning process underlying ones' intuitive actions.

Selected rationality emphasizes the process of selection of decision procedures through survival or growth; choice is dominated by rules that have survived and evolved. Contrary to the decision analysis dictum of rewarding good decisions rather than good outcomes, this type of rationality stems from a system that rewards good outcomes (an individual who takes a poor risk would be rewarded if he happens to have a lucky outcome). In this case accountability appears to be a scarce resource; it is often much easier to measure the outcomes of a decision than to review the process by which the decision was made.

Posterior rationality emphasizes the discovery of intentions as an interpretation of action rather than as a prior position; thus, decisions are antecedent to goals and define preferences rather than follow them. The scarce resource here seems to be foresight or clear thinking, since this type of rationality implies taking actions before fully evaluating the consequences.

In principle, decision aids can reduce the bounds inherent in both selected and posterior rationality. However, it seems clear that no aid can hope to completely eliminate these bounds. For example, in the context of selected rationality, a decision analytic aid could provide a consistent framework for explaining the elements that went into making a given decision, and help a user appraise the quality of the decision process. On the other hand, it is not feasible to fully analyze every decision, even if every decision is made using the tools of decision analysis. Therefore, it seems inevitable that some form of selected rationality must remain in spite of the application of an analytic aid. Similarly, although an analytic aid can make it easier to choose actions that are based on one's preferences, it is likley that some sort of posterior rationality will always remain, particularly for minor decisions.

The case for decision aiding is most controversial in the context of adaptive rationality. It is this type of rationality that critics like Tribe (19), and Dreyfus and Dreyfus (2), defend to the exclusion of formal methods. Hoos (3) also discusses this problem. They feel that the quantification process required by many analytical aids can be inherently bad. Even if a formal analysis could lower the bounds inherent in adaptive rationality (a point they do not really concede), they do not believe that lowering the bounds is necessarily a good thing. They argue that the "intangible" elements that can't be quantified are often the most important parts of the problem, and that the process of formalizing a decision-maker's intuitive process tends to ignore these factors.

If we take the traditional view of either accepting or rejecting the conclusions derived from an analytical decision aid, then the critics' points seem well taken. It is certainly not reasonable to assert that adaptive rationality is always wrong or naive; most of us know persons that seem to be naturally good decision-makers. There is no logic that guarantees that an aid will be better than the decision-maker's intuition every time. On the other hand, most would agree that some of the benefits of decision aiding — the ability to explore the logic of decisions, the ability to handle uncertainty and complexity, the expanded capability for communicating the results to others — are very attractive. The challenge is how to take advantage of the benefits of decision aids without throwing away the rich body of intuitive knowledge and judgement that often cannot be formalized and quantified.

To summarize, there are a number of alternate forms of descriptive rationality. Each of these rationalities can be interpreted as the result of some bound or scarce resource constraint on the ability to formulate, think about, and solve decision problems. Formal decision-aiding methods have the potential to overcome these bounds in many cases. However, when decisions aids are used to the exclusion of unaided judgement and intuition, there is a valid question as to whether the aids do more harm than good.

The view that a conclusion derived from a decision aid must be completely accepted or rejected is perhaps the result of the cultural assimilation of the "scientific method" into our thinking. In the scientific method, a hypothesis is proposed and then either accepted or rejected on the basis of some experimental evidence. This type of thinking, which has served so well in the field of scientific discovery, is wholly inappropriate in the field of decision-making.

A broader and more useful view of a decision aid is that it provides a decision-maker with additional data on which to base his choices. The result of decision-aiding is information for the decision-maker, not a substitute for intuition, good judgement or clear thinking. Viewing analytical aids as complements to, rather than substitutes for, unaided intuition frees us from the unnecessarily rigid vision of formal analysis as a competitor of the intuitive process.

In science the only place for intuition is the design of an experiment; in the unavoidably personalistic realm of decision-making, intuition must play a central role. When faced with a one-of-a-kind decision problem, a decision-maker can do no better than act based on what he has (alternatives and resources), what he wants (values and preferences), and what he knows (information). Aids can help provide information and clarify choices, but it is unrealistic and unnecessary to expect them to wholly replace the decision-maker.

In effect we are claiming that an aid should be viewed as a special kind of expert. No one would ever argue that all expert opinions should be believed, or that all expert advice should be followed to the letter. It makes just as little sense to treat the conclusions derived from an analysis as inviolate. However, it is equally shortsighted to completely ignore an aid if it can provide insights, but fails to provide a perfectly credible answer. The information produced by the aid, just like the advice of the expert, must be tempered by the decision-maker's feelings about the credibility of the analysis that produced it. The formal development of these intuitive ideas has been started by Morris (9), and Nickerson and Boyd (12).

In summary, we have seen that the actual process by which humans make decisions is quite complex. We have attempted to argue that because of, not in spite of, the limited rationality with which the average person can hope to apply to a decision, the prospects for helpful decision aiding are good. On the other hand, the naive use of an analytical aid may well do more harm than good, and its user must clearly understand its proper role in order to ensure that it makes a positive contribution.

The following chapter will explore in more detail the idea of using an aid as an expert and will use this approach to derive a measure of the aid's credibility.

8. THE VALUE AND CREDIBILITY OF DECISION AIDS

This chapter describes three different, but related, analytic procedures for determining the value and credibility of decision aids. Different analytic approaches are used for investigating the characteristics of different types of decision aids. There is a spectrum of decision aids ranging from simple information processing devices to sophisticated systems for formulating and analyzing decisions. The first two procedures described in this chapter deal with classes of aids at either end of this spectrum, and the third procedure is a general conceptual approach to determining the value of all types of aids. Some of these procedures are more fully developed than others, and some potentially have broader application. With further research we expect each of these approaches to become part of a general theory of the credibility and value of aids.

The first section of this chapter addresses the issue of measuring the value of a conceptually simple decision aid, such as one that displays the location of objects detected by radar. Although such devices are not normally considered decision aids, they represent one end of the spectrum of the information processing functions described in Chapter 3. For these physical measurement devices, their value may be related directly to the measurement error.

The second section, which describes the majority of our work on credibility, examines the class of sophisticated decision aids that not only provide inputs to a decision-maker, but actually help in the process of formulating and making a decision. For such aids the issue of credibility is very complex, but it is an important factor in determining an aid's value. An aid that a decision-maker does not believe has little value, no matter how sophisticated it may be. It is not obvious how to measure the credibility of such aids, and even less obvious how the value of the aid is related to the degree of credibility. However by carefully defining the meaning of credibility in this context, we are able to specify a step-by-step procedure for determining the credibility of an aid for decision structuring.

The third section presents a general conceptual framework for evaluating aids. We define an analytical construct called a "stochastic decision tree" which appears to be a completely general way to think about the value of an aid in a variety of decision situations. Although this work does not bear directly on the issues of credibility, it is our hope that stochastic decision trees can be used to unify the concepts developed in the first two sections.

The Value of Decision Aids for Assessing or Measuring Parameters

This section presents an analytic framework for evaluating the benefits of a class of simple decision aids. We examine decision aids that provide information to the decision-maker rather than help him in the logical process of making his decision. While the approach is compatible in concept with the decision tree methods described in the next section, we are able to obtain stronger results by narrowing our focus. The hope is that this specialized methodology can be extended to other aids, and, with further research, can be merged with the other concepts described in this chapter to form a general procedure for evaluating decision aids.

We assume that the information produced by the aid allows the decision-maker to make a more accurate estimate of a variable or factor that is important to the decision he is facing. We characterize each potential decision situation for which the aid is relevant by three measures: the importance of the situation, the importance of having an accurate estimate, and the decision-maker's prior uncertainty about the uncertain factor. The methodology quantifies the value of an aid in terms of these three measures.

By then specifying the frequency with which different types of decision situations occur, we can quantify the overall value of an aid. These ideas are explained in more detail below.

Framework

Assume for a given decision situation that there is a true state of the world. The aid provides information that leads to an estimate of the true state. For example, consider the problem of locating an aircraft. The true state of the world is the actual location of the aircraft. A decision aid that displays the location of objects detected by radar helps to pinpoint the aircraft's location.

On the basis of the information supplied by the aid, the decision-maker makes a new estimate of the actual state of the world. This estimate may or may not be equal to the estimate the aid provides. In many cases the decision-maker may choose to incorporate his own intuition about the true state of world, as well as the information provided by the aid.

We shall posit that in a given decision situation there is an underlying cost function that describes the cost of an inaccurate estimate. The simplest such function would assign a fixed cost to incorrectly estimating the true state of the world (as, for example, in a "hit or miss" situation). A more general cost function would describe how the cost increases as the estimate diverges from the true state of the world.

The final ingredient in our formulation is an assessment of the decision-maker's uncertainty about the true state of the world prior to using the decision aid. If the decision-maker were absolutely certain about the state of the world, then the value of the aid would be zero. The more uncertain the decision-maker is about the true state of the world, the higher the value of a decision aid that provides information about the true state.

Thus, we describe a given decision situation in terms of three characteristics.

- The magnitude of the decision problem (i.e., a measure of its importance).
- A measure of the importance of accurately estimating the true state of the world in a given situation.
- The decision-maker's uncertainty about the true state of the world in the given situation.

For each specific setting of the above characteristics we can compute the value of an aid by looking at the value of the decision situation with and without using the aid.

Any decision aid probably will be used in a sequence of different decision situations. By characterizing the frequency with which each decision situation occurs, we can calculate the total value of the aid by multiplying the value in each situation by the number of times the situation occurred, and aggregating this value over all possible situations.

The above ideas provide what appears to be a natural framework for quantifying the evaluation of this class of decision aids. We have experimented with several mathematical models that attempt to quantify and build a rigorous analytic foundation for the framework described above. For brevity, we omit the detailed mathematical description here. The basic result is that it appears that we can quantify the characteristics of any given decision situation in terms that a decision-maker can understand and assess. With further research we believe that it will be possible to apply the framework to the evaluation of real decision aids.

Credibility of Decision Tree Aids

This section reports the results of our research on decision aid credibility. The research has three main objectives:

i. Define precisely the concept of credibility

ii. Determine a quantitative measure of credibility that can be used to evaluate different decision aids

iii. Develop an operational procedure for assisting decision-makers in utilizing decision aids.

There are many different types of aids that could help a decision-maker structure and make a decision. To bound the problem we focused on one representative aid whose purpose is to help the decision-maker structure a decision tree, and then evaluate the decision tree to recommend the best alternative. The aid would interact with the decision-maker to elicit both the structure of the decision and the probabilities and values that characterize the tree. A sophisticated version would provide additional assistance in helping the decision-maker provide these inputs (through a structured inquiry system or possibly through some additional modeling).

As discussed in Chapter 6, a large body of descriptive decision-making literature supports the conclusion that the credibility of a decision aid (or analysis, or model) influences the way it will be used. This existing work indicates the importance of the issue. However, there is little theoretical work on the subject of how the credibility of an aid should influence the way it is used. The closest related work has been in the area of model evaluation. This work is relevant for the decision aids we are examining, aids that produce models to assist in the decision-making process. We have included a list of references at the end of this report.

Approach

Our approach is best understood in the context of the traditional role of decision models. We begin with a brief review of the classical decision analysis paradigm.

In the traditional view, there are two solutions to a given decision problem. These are depicted in Figure 9. The first solution is the decision-maker's intuitive decision; i.e., the decision he would make without any formal analysis or model. The second solution is the model solution. We assume that a decision tree model has been constructed that structures the alternatives and outcomes and, based on a set of probability and value judgements, determines an "optimal" decision. This decision is called the "model decision".

In the traditional paradigm, the decision-maker must decide to accept or reject the model. If he rejects the model, he makes the intuitive decision; if he accepts the model, he chooses the model decision. There are at least two problems with this. First, there are no explicit criteria for determining whether to accept or reject the model. The decision-maker is often placed in the position of having to judge a model whose technical details he doesn't really understand. Second, and perhaps more important, in rejecting either the model or his intuition, the decision-maker is throwing away potentially useful information. The danger of neglecting intuition is the danger of neglecting important insights just because they are hard to model; the danger of neglecting the model is the danger of neglecting valid information and logic just because it is hard to understand.

A more general approach that we shall take is to view the model as providing new information about the value of the decision alternatives. In this view, the question of accepting or rejecting the model is meaningless. The model simply provides a source of information that the decision-maker can use

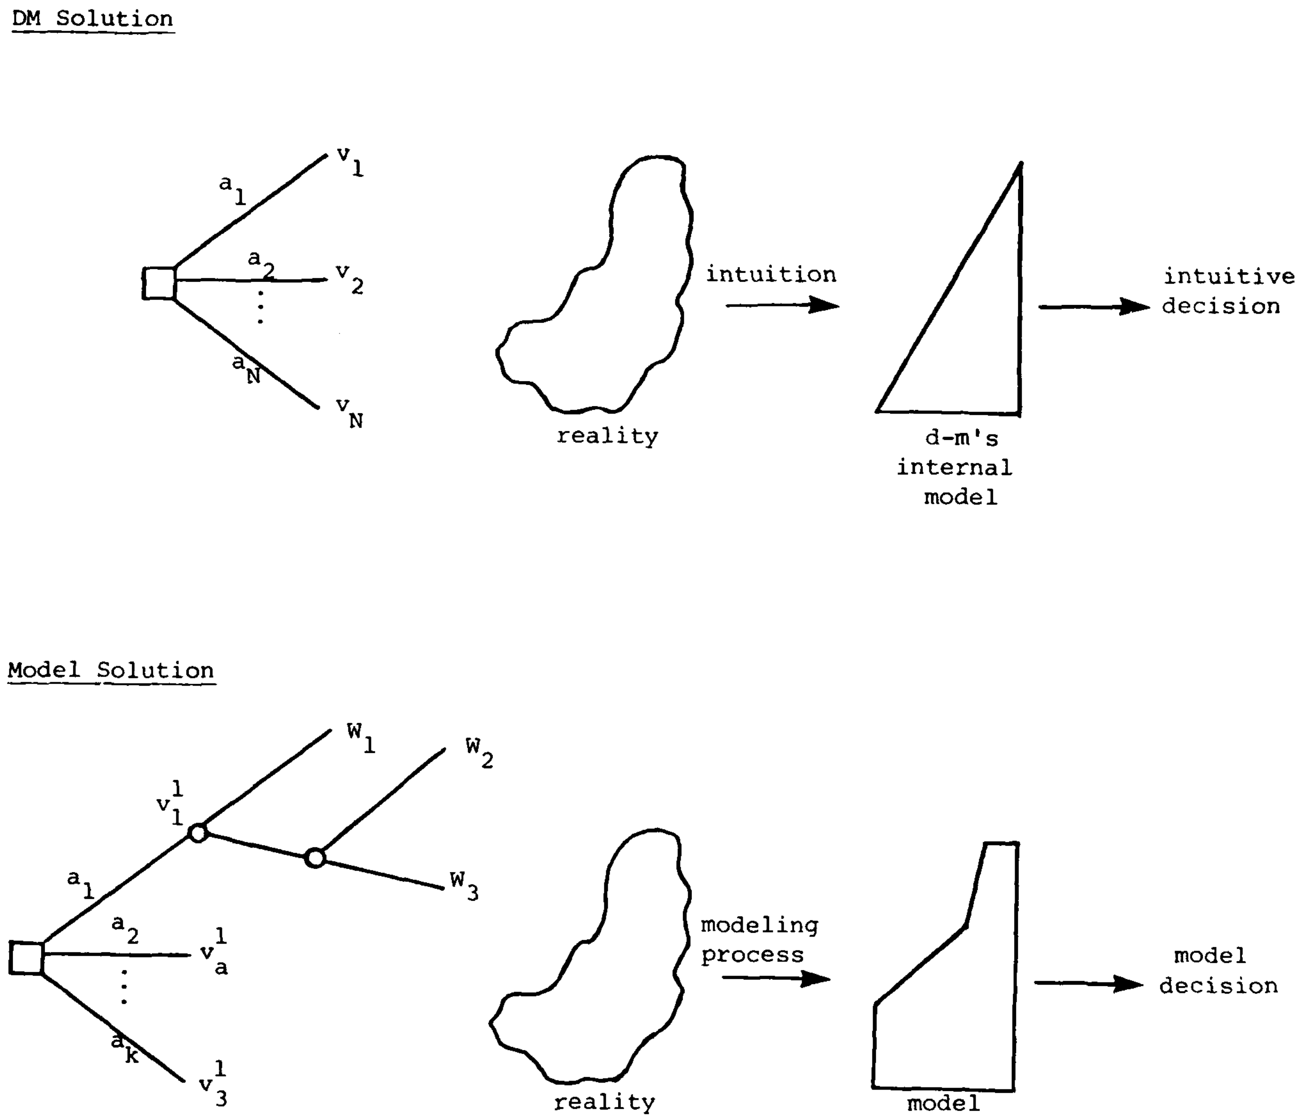

THE CLASSICAL PARADIGM

Figure 9

to update his intuitive judgement about the value of each alternative. The updated values, based on the information supplied by the model, are then used to select the best alternative. This broader formulation treats the model as a complement rather than a substitute for intuition.

A feature of the analysis is that it provides a framework for determining the value of further analysis. For decision aids that help the decision-maker structure and think about a problem, this is especially important. For these aids there is a continuous decision of whether to stop using the aid and make the final choice, or continue using the aid to structure the problem further.

Credibility of a Model

It is easy to speak and think of credibility as being a physical attribute of a model. However, this is not a very useful way of thinking since two people may have very different feelings about the credibility of the same model. A more productive view is that the credibility of a model is related to how accurately it represents one's perceptions of reality. Since representation is by nature a subjective process, so must be its evaluation.

Another fallacy that goes to the heart of evaluating the credibility of a model is that there is only one correct model. It is easy to fall into the conceptual trap of believing that if we just had enough time we could achieve a perfect representation of reality. The logical flaw in this is that no matter how detailed a model is, we could always add additional structure. In fact, the best we can do is accurately represent our perception of reality, which is always limited. The following simple example will amplify these ideas and provide a useful reference for the discussion that follows.

Example: The Cuban Missile Crisis

Consider the simple decision tree illustrated in Figure 10 that depicts President Kennedy's decision in the Cuban missile crisis. One of his options was to blockade Russian ships. If he chose a blockade, he had to be concerned with the possibility of precipitating a nuclear war. Let us assume there was some probability p of a war occurring. In this simple model, there are three possible outcomes labeled w_1, w_2, and w_3 which are the values the President would attach to each outcome.

Figure 10 also shows a more detailed model of the same situation. Model II represents more of the problem structure, including the possibility that the Russians might try to run the blockade with some probability q, resulting in two different conditional probabilities of nuclear war, r and s. The model also includes the option of economic reprisal in the case of no blockade. In this augmented decision tree there are six basic outcomes on which the President would have to assign values u_1 through u_6.

Most observers would agree that Model II is in some sense more credible than Model I. It is more detailed and includes more of the actual problem structure. However, it is easy to show using elementary probability theory that Model II could be reduced to a form exactly equivalent to Model I (for example, the probability of nuclear war would be $(qr+(1-q)s)$. The only possible difference would be in the numerical values of the probabilities and values.

Since the models are logically consistent, it makes no sense to think of either structure as being more or less correct. It is also apparent that an assessment of credibility must rely in part on information not available by measuring characteristics of the model such as number of variables, level of detail, etc. For instance, it is easy to construct examples of more detailed decision models of the Cuban missile crisis that are consistent but less credible than those in Figure 10.

AD-A086 716 APPLIED DECISION ANALYSIS INC MENLO PARK CA F/G 17/2
ANALYTIC PROCEDURES FOR DESIGNING AND EVALUATING DECISION AIDS.(U)
APR 80 A C MILLER, P A MORRIS, R D SMALLWOOD N00014-78-C-0240
UNCLASSIFIED ADA-2017 NL

2 of 2
AD
A086716

END
DATE
FILMED
8-80
DTIC

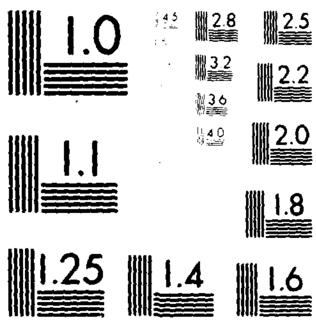

MICROCOPY RESOLUTION TEST CHART

Two Decision Models:

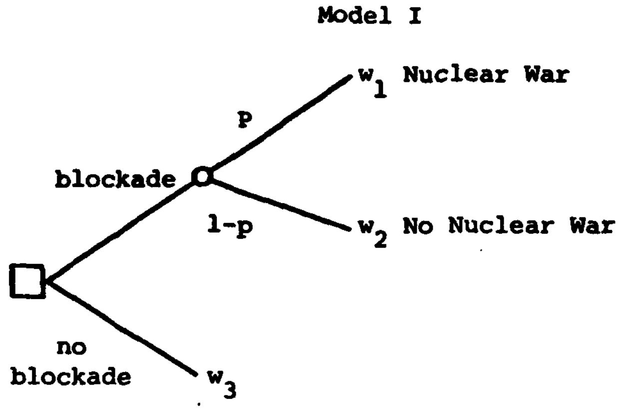

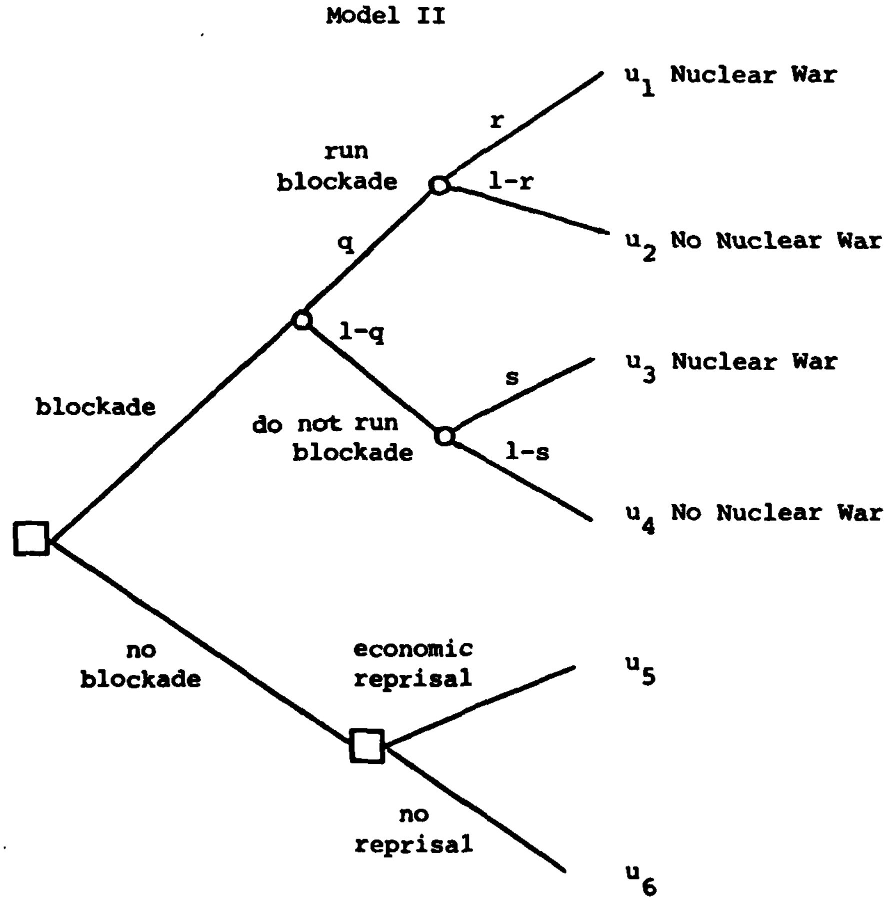

EXAMPLE - THE CUBAN MISSILE CRISIS

Figure 10

Perfect Credibility

A more relevant question to ask about the two models is how might the analysis change with additional modeling or information gathering. Consider two cases, corresponding to different feelings a decision-maker might have about a given model:

Case 1: There is a good chance that further modeling will significantly change his assessment of the value of each alternative.

Case 2: There is no chance that further modeling or information would lead to any new beliefs.

Most would agree that the credibility of the model in Case 2 is higher than in Case 1. Intuitively, if we know that our modeling assumptions are absolutely valid in the sense that no information we might observe, or modeling we might perform, prior to the decision would change the answer, then the model is as credible as possible. Conversely, if our assumptions are likely to change with more modeling or feasible information gathering, then the current model is not completely credible.

Consider again Model I, the simple model of the Cuban decision. Why is this not a very credible model? One reason is that the model, in effect, forces the decision-maker to specify probabilities of very complex outcomes. Thus, we might expect the decision-maker to feel that the probability assessment in Model I is in some sense relatively inaccurate, and likely to change with more thinking and modeling. Similarly, if we ask him to assign a value to the consequences of no blockade, he might be quite uncomfortable doing so since it depends on a number of unstated factors (some of which are included in Model II). Considering these factors explicitly may well affect the assessed value of this alternative.

Model II on the other hand decomposes the problem into more basic elements. Presumably, the decision-maker would be more comfortable assessing the more basic elements and having the model perform logical computations than he would be if he had to keep track of everything in his head. We would expect him to be more confident that further modeling would not change the results.

In summary, we have found that the notion of model credibility is related to the decision-maker's feelings about how future modeling and information could affect the model results. Furthermore, the stability of the model results is a function of the stability of the model inputs and assumptions. If all the assumptions were invariant to future information or modeling, then the model would be perfectly credible. We are thus led to a formal definition of a "perfectly credible model". This definition will be useful both as a conceptual thinking tool and for quantitative analysis.

<u>definition:</u> A perfectly credible model is a model whose results are estimated by the decision-maker to be invariant to additional feasible information gathering or modeling.

It is important to distinguish the notion of a "perfectly credible model" from the notion of a "perfect model" as proposed by Nickerson and Boyd (12). They define a perfect model as one that predicts with certainty the value of any alternative. There are two reasons why this is unacceptable as a definition of a perfectly credible model. First, such a model is unattainable in virtually all real decision problems. Second, many non-deterministic models are completely credible in the usual sense of the word. A classic example is modeling a sequence of coin flips as a Bernoulli process with probability 0.5 of heads or tails. Intuitively, the Bernoulli model is completely credible because we believe no amount of information or modeling would change our conclusion that the probability is 0.5. Although we could imagine extreme situations in which we could make more refined observations (thumb velocity, height from the floor, resiliency of the coin, etc.), in most cases it is not possible to make such observations before the decision.

Our definition of a perfectly credible model takes into consideration the amount of information feasible to observe before the decision. This of course also limits the scope of the modeling that may be performed before the decision.

Sources of Model Credibility

It is useful to think about decision-tree models such as Model I and Model II as black boxes that have certain inputs and outputs. The inputs to each model are the probabilities describing the chances of uncertain outcomes, and the values describing the worth of those outcomes. The decision analysis methodology provides an explicit logical way to determine the best alternative given these inputs. This is the alternative with the highest expected value. (For a risk averse decision-maker, these values will be utilities which reflect a desire to minimize the uncertainty in the outcomes.)

There are really three basic types of inputs to the description of the decision problem. In addition to the probabilities and values, there is the structure of the decision tree itself; i.e., the alternatives and events described by the tree. A useful way to view both the probability and value assignments is to regard them as summaries of situations that could always be structured in more detail. For example, the assignment of a value to the no-blockade option in Model I must take into account, at least implicitly, other possible actions that might be taken. The issue is whether to include this structure explicitly in the formal model, or implicitly in the assessment of the value.

It is always possible to interpret an event probability as an estimate of the unknown frequency of occurrence of similar events. If we are certain about the frequency, then the probability assignment won't change with further information or modeling. However, if the frequency is unknown, then additional modeling or information will potentially alter the probability. We are thus led to a natural definition of a "perfectly credible probability" as one on which there is no uncertainty concerning the underlying frequency. The probability assignment of 0.5 to a fair coin is perfectly credible because the frequency is virtually known and no feasible modeling or information would change it. On the other hand, a probability assignment of 0.5 to the event that our next president will be Republican is not perfectly credible because additional modeling or information gathering could easily change it. In most circumstances a perfectly credible probability will be unattainable. However, the definition is useful because it represents the best we can do.

Similarly, we define a "perfectly credible value" as one that does not change with further modeling or information. Conceptually, we could obtain perfectly credible values by building a sufficiently detailed decision tree so that the outcomes are unambiguously defined and do not include implicit uncertain events or unstructured alternatives.

An expected value calculated by rolling back a part of a decision tree will be perfectly credible if and only if all the probabilities and values in that portion of the tree are themselves credible. Thus, a basic consistency condition is that a model is perfectly credible if and only if all its parameters are.

A useful analogy to the process of specifying probabilities and values is the physical measurement process. Just as there are measurement errors when making measurements of physical phenomena, so there will be assessment errors in specifying probabilities and values. Thus, we presume that both the intuitive value and the model value are, in general, not equal to the perfectly credible value; rather they both are estimates of this value that contain, in general, some error.

The decision-maker's prior uncertainty about the perfectly credible value of each alternative may be assessed directly as a probability distribution. This assessment, which we shall discuss later, is a measure of the decision-maker's confidence in his own intuition. We call it the "intuitive assessment".

The uncertainty in the perfectly credible value implied by the model may be calculated from an assessment of the uncertainty in the errors of the inputs. The decision-maker assesses the uncertainty in each parameter and then the model is used to calculate how the uncertainty propagates through the tree. (This calculation is a straightforward application of probability theory.) We call the resulting probability distribution the "model-based assessment." The model-based assessment is a measure of the decision-maker's feelings about the accuracy of the model. If the model were a human expert, the model-based assessment would be a measure of the expert's confidence in his advice.

It is useful and intuitive to think of the range of uncertainty represented by the model-based assessment as a measure of the credibility of the model. The more certain the model-based assessment, the more credible the model. The following definition of model credibility makes this notion precise.

<u>definition:</u> The <u>credibility</u> of a model is equal to the precision of its probability distribution on the perfectly credible value.

The precision of a probability distribution is a well-defined statistical measure equal to the reciprocal of the variance. Infinite precision corresponds to complete certainty. Thus, a perfectly credible model is one whose credibility is infinity.

The evaluation of a given model depends on the relative credibility of the model and the decision-maker's intuitive assessment (his intuitive model). The next section presents a method for performing this evaluation.

A General Result

We are now in a position to determine how the decision-maker should update his intuitive assessment, based on reception of the model-based assessment. This is the central analytic result of our research.

We regard the model result as information. This information may be thought of as the result of an experiment, the modeling process. Viewed in this way, the updating problem is a classical inference problem, and Bayes' theorem may be applied to determine the posterior probability distribution on the perfectly credible value of each alternative, based on observation of the model-based assessment.

The posterior distribution depends on the likelihood function, which is an assessment of the relative likelihood of every possible model the aid might produce. Unfortunately, although it is possible in theory to assess the likelihood function directly, for most reasonably sophisticated decision aids, direct assessment is practically impossible.

One way to simplify the assessment is to approximate the likelihood function. A result of our research is that the likelihood function can be obtained directly from an assessment of the credibility of the model if the shape of the likelihood function is similar to (or can be approximated by) that of a Normal distribution, and if two additional behavioral assumptions are satisfied.

The first behavioral assumption is that, if the decision-maker has no prior knowledge, or has done very little thinking about the decision problem, then he will adopt the state of information represented by the model as his own. This assumption would be accurate if the decision-maker does not feel that the model has any systematic bias. In fact, a good decision aid will be designed to help the decision-maker structure his decision problem to eliminate just such systematic biases.

The second behavioral assumption is that the result of the formal modeling process may be viewed as the outcome of an independent experiment. This means practically that our results will apply only to reasonably extensive modeling efforts that extend the decision-maker's thinking in a substantial way. Decision aids that result in very simple models (as in, say, a short-term crisis situation) would likely produce results highly dependent on the results of the decision-maker's intuitive model. Although we have some preliminary ideas or how to characterize such dependence, we have not addressed the subject in our current research.

For brevity, we shall skip the mathematical details here and show the implications of these assumptions. Let v_d and c_d be the expected value and credibility (the precision) of the decision-maker's intuitive assessment of a given alternative, and let v_m and c_m be the expected value and the credibility assigned by the model. We are interested in the updated value and credibility, based on observance of the model result. The result is that the new parameters are related to the decision-maker's prior evaluation and the model's assignment in a very simple way:

$$\text{updated value} = \frac{c_m v_m + c_d v_d}{c_m + c_d}$$

$$\text{updated credibility} = c_m + c_d$$

Both the model-based value v_m and the intuitive value v_d are weighted by their relative credibilities.

The above result allows us to determine the posterior distribution for the perfectly credible value for any decision alternative from a set of relatively simple assessments. The evaluation is based on an explicit measurement of the credibility of the model. The credibility measurement is derived from a simple set of assessments. We discuss the assessment task below, and conclude with a step-by-step summary of the methodology.

The Assessment Task

Apart from the task of building the model, the decision-maker must make three types of assessments in order to encode the credibility of a decision tree model. First, he assesses the possible variation in the intuitive values he assigns to each alternative before using the aid. Then he encodes his judgement about the possible variation in the probability inputs to the model, and his judgement about the possible variation in the values at the tips of the tree.

The possible variation in each parameter can be summarized by the credibility or precision of the probability distribution on that parameter. The easiest and most direct way to estimate the precision is to have the decision-maker provide a 70 percent probability range such that there is equal probability that the revealed value will fall on either side of the nominal value. Since the standard deviation of a probability distribution defines a range having approximately 68 percent probability (exactly 68.2 percent for the Normal distribution), we may use the 70 percent range as an estimate of the standard deviation. The precision, or credibility, is then calculated directly as the reciprocal of the square of the width of the range.

To summarize, the crediblity assessment task consists of specifying a single 70 percent range for each parameter that is not perfectly credible.

The Step-by-Step Procedure

We now summarize the methodology in terms of a procedure for utilizing a model produced by a decision aid:

Step 1 - Encode Intuitive Judgement

The decision-maker attaches a nominal value and a range to each alternative. The result is the best decision based on intuition and a measure of the credibility the decision-maker attaches to his intuition.

Step 2 - Construct the Model

Using the decision aid, the decision-maker constructs a decision tree and assigns nominal probabilities and values.

Step 3 - Encode Parameter Credibility

The decision-maker assigns a 70 percent range to each probability and value in the tree.

Step 4 - Calculate the Model Credibility

The model credibility is calculated from the parameter credibilities. As discussed earlier, this calculation is done using the model itself.

Step 5 - Calculate the Updated Expected Value of Each Alternative

The updated value is computed using the relative credibilities of the model and the decision-maker's intuition.

Step 6 - Calculate the Credibility of the Updated Model

The posterior credibility is computed as the sum of the intuitive and model credibilities. This is useful for evaluating further modeling efforts.

It is important to note that the decision-maker is only involved in the first three steps of the methodology. Steps 4, 5, and 6 are all logical operations that could be performed by the aid itself. Thus, the decision aid could be designed so that it assists in its own evaluation.

An important benefit of augmenting the aid to assist the decision-maker assess its credibility is that we can use the posterior credibility as a gauge of the value of refining the model. The new paradigm, which treats the model result as information, allows us to utilize the powerful results of information-value theory. In particular, it is possible to calculate the value of a perfectly credible model. This provides an upper bound on the possible benefits of using an aid that helps the decision-maker produce a decision model.

The ideas described above appear to be promising in terms of practical application to decision tree aids. Furthermore, the ideas can be extended to a broader class of aids (e.g., aids of known structure with a finite number of inputs). The main issue that needs additional attention is that of the dependence between the modeling process and the process the decision-maker uses when making unaided decisions. The results stated here assume that these processes are independent. Further research is necessary to explore more thoroughly the issues of probabilistic dependence, both in parameter estimation and between the formal model and the intuitive model.

Stochastic Decision Trees

One important dimension for classifying decision problems is the degree to which they are similar to other decisions faced by the same decision-maker. At one end of the spectrum we might have decisions that are virtually identical to a number of other decisions, while at the other end of the spectrum we have decisions that are completely different in structure and content. For example, the decisions associated with launching successive aircraft from a carrier are very similar. On the other hand, situations such as the Cuban missile confrontation represent decisions that a decision-maker might face only once in a lifetime.

Somewhere between these two extremes are decision situations that are similar in structure but have differences in degree that can be represented by a set of parameters. For example, the detection of an unidentified aircraft approaching a task force creates a decision situation that can be typified by several parameters such as weather conditions, readiness of the fleet, trajectory of the aircraft, and the consequences of failing to identify correctly and intercept the aircraft.

This kind of situation can be characterized by a decision tree that identifies the initial actions, potential responses, secondary actions and ultimate resolution of the encounter. The part of the problem that will change from situation to situation is the parameters of the decision tree (i.e., the probabilities and values placed on the outcomes). In principal, one could imagine drawing a decision tree for each generic type of situation and then characterizing the possible situations likely to be faced by the decision-maker by a joint probability distribution over the probabilities and values in the decision tree. We have termed such an entity—the decision tree structure and the joint probability distribution over its parameters—a "stochastic decision tree". One can imagine constructing such a stochastic decision tree for each of the major decision situations faced by a decision-maker, or at least those decision situations that have similar structures.

From a practical point of view it is not likely that the joint probability distribution required in the specification of the stochastic decision tree could be constructed directly from known data or experience. Instead, we could build models that characterize various scenarios likely to be faced by the decision-maker and then use these models to generate the required joint probability distributions.

Once a stochastic decision tree has been specified it is possible to calculate the potential value of a decision aid to the class of decisions it represents. For example, if one were considering a decision aid that supplies additional information about one of the state variables, then it is possible to calculate the value of that additional information to the class of decisions represented by the stochastic decision . This calculation can be carried out in one of two ways: An analytical expression can be derived for the value of each alternative action as a function of the probability for the specific variable under consideration. With this calculation it is possible to compute the expected value of additional information about the unknown variable. If this calculation proves to be too difficult, the value of the additional information about the stochastic decision tree can be calculated using a Monte Carlo technique. In this technique, random samples are chosen from the joint probability distribution for the parameters of the stochastic decision tree and the value of the additional information for this particular setting of the parameters is calculated. Successive samples then yield the distribution over the value of information.

Although stochastic decision trees seem like a natural way to describe decision situations in which the structure is constant but the parameters are not, more research is required on this topic. This research should be directed toward two important topics: a) validation of the general approach as a useful descriptor of military decision situations, and b) development of methodologies for using the concept to calculate the value of decision aids under the assumptions of normative decision-making.

REFERENCES

1. B. De Finetti, "Foresight: Its Logical Laws, Its Subjective Sources", Studies in Subjective Probability, Kyburg and Smokler, Ed., Wiley, New York, 1963.

2. H. Dreyfus and S. Dreyfus, "Uses and Abuses of Multi-Attribute and Multi-Aspect Models of Decision-Making", Unpublished Report.

3. I. Hoos, Systems Analysis in Public Policy, University of California Press, Berkeley, 1972.

4. R.A. Howard, "Foundations of Decision Analysis," IEEE Transactions on Systems Science and Cybernetics, Vol. SSC-4, No. 3, September 1968).

5. J.G. March, "Bounded Rationality, Ambiguity, and the Engineering of Choice", Bell Journal of Economics, 1979.

6. J.E. Matheson, "The Economic Value of Analysis and Computation", The IEEE Transactions on Systems Science and Cybernetics, Vol. SSC-4, No. 3, 1968.

7. A. Miller, T. Rice, M. Metcalfe, "An Analytic Characterization of Navy Command and Control Decisions," March 1979, ADA Final Report, ONR Contract Number N00014-78-C-0240.

8. A.C. Miller, J.V. Rowney, J.R. Payne, "The Naval Task Force Decision Environment," September 1974, SRI Technical Report NWRC-TR-8, ONR Contract Number N00014-74-C-0357.

9. P.A. Morris, Bayesian Expert Resolution, Ph.D. Dissertation, Stanford University, 1971.

10. P.A. Morris, "Combining Expert Judgments: A Bayesian Approach", Management Science, Vol. 23, 1977.

11. P.A. Morris, "Decision Analysis Expert Use", Management Science, Vol. 20, 1974.

12. R.C. Nickerson and D.W. Boyd, "The Use and Value of Models", to appear in Operations Research, 1979.

13. R. Radner, "Satisficing", Journal of Mathematical Economics, Vol. 2, 1979.

14. T.R. Rice, "The Value of Data For a Quadratic Decision Problem", Management Science, Vol. 24, 1978.

15. H.A. Simon, "The Behavioral Model of Rational Choice", Quarterly Journal of Economics, Vol. 69, 1955.

16. H.A. Simon, "On How to Decide What to Do", Bell Journal of Economics, 1979.

17. R.D. Smallwood, "A Decision Analysis of Model Selection", The IEEE Transactions on Systems Science and Cybernetics, Vol. SSC-4, No. 3, 1968.

18. S.N. Tani, "A Perspective on Modeling in Decision Analysis, Management Science, Vol. 24, 1978.

19. L. Tribe, "Policy Science: Analysis or Ideology?", Philosophy and Public Affairs, 1972.

DISTRIBUTION LIST

Director, Engineering Psychology
Programs (Code 455)
Office of Naval Research
800 North Quincy Street
Arlington, VA 22217 (5 cys)

Defense Documentation Center
Cameron Station, Bldg. 5
Alexandria, VA 22314 (12 cys)

Dr. Craig Fields
Director, Cybernetics Technology
Office
Defence Advanced Research
Projects Agency
1400 Wilson Boulevard
Arlington, VA 22209

Naval Analysis Programs
Code 431
Office of Naval Research
800 North Quincy Street
Arlington, VA 22217

Statistics and Probability
Program
Code 436
Office of Naval Research
800 North Quincy Street
Arlington, VA 22217

Information Systems Program
Code 437
Office of Naval Research
800 North Quincy Street
Arlington, VA 22217

Dr. Jesse Orlansky
Institute for Defense Analyses
400 Army-Navy Drive
Arlington, VA 22202

Commanding Officer
ONR Branch Office
ATTN: Dr. Charles Davis
536 South Clark Street
Chicago, IL 60605

Commanding Officer
ONR Branch Office
ATTN: Dr. E. Gloye
1030 East Green Street
Pasadena, CA 91106

Analysis and Support Division
Code 230
Office of Naval Research
800 North Quincy Street
Arlington, VA 22217

Operations Research Program
Code 434
Office of Naval Research
800 North Quincy Street
Arlington, VA 22217

Commanding Officer
ONR Branch Office
ATTN: Dr. J. Lester
Bldg. 114, Section D
666 Summer Street
Boston, MA 02210

Dr. James McGrath
Navy Personnel Research and
Development Center
San Diego, CA 92152

Dr. Donald A. Topmiller
Chief, Systems Effectiveness
Branch
Human Engineering Division
Wright Patterson AFB, OH 45433

Dr. Julie Hopson
Human Factors Engineering
Division
Code 604
Naval Air Development Center
Warminster, PA 18974

Dr. Edgar Johnson
Organizations & Systems Research
Laboratory
U.S. Army Research Laboratory
5001 Eisenhower Avenue
Alexandria, VA 22333

Dr. S.D. Epstein
Analytics
2500 Maryland Road
Willow Grove, PA 19090

Dr. Amos Freedy
Perceptronics, Inc.
6271 Variel Avenue
Woodland Hills, CA 91364

Dr. G. Hurst
University of Pennsylvania
Wharton School
Philadelphia, PA 19174

Dr. Miley Merkhofer
Stanford Research Institute
Decision Analysis Group
333 Ravenswood Avenue
Menlo Park, CA 94025

Dr. Kenneth Gardner
Applied Psychology Unit
Admiralty Marine Technology
Establishment
Teddington, Middlesex TW11 OLN
ENGLAND

Dr. C. Kelly
Decisions and Designs, Inc.
8400 Westpark Drive, Suite 600
P.O. Box 907
McLean, VA 22101

Dr. Arthur Siegel
Applied Psychological Services
Science Center
404 E. Lancaster Street
Wayne, PA 19087

Mr. David Walsh
Integrated Sciences Corporation
1640 Fifth Street
Santa Monica,CA 90401

I.R. Mirman
Assistant for Special Projects
HQ AFSC-DL
Andrews AFB, MD 20334

Mr. Joseph Wohl
MITRE Corporation
Box 208
Bedford, MA 01730

Dr. Rex Brown
Decision Science Consortium
Suite 421
7700 Leesburg Pike
Falls Church, VA 22043

Mr. Philip Andrews
Naval Sea Systems Command
NAVSEA 0341
Washington, D.C. 20362

Dr. Robert Kolb
Naval Oceans Systems Center
271 Catalina Boulevard
San Diego, CA 92106

Mr. Leslie Innes
Human Factors Wing
Defense & Civil Institute
of Environmental Medicine
P.O. Box 2000
Downsview, Ontario
M3M3BQ
CANADA

Dr. Albert Colella
Combat Control Systems
Naval Underwater Systems Center
Newport, RI